System Administrator's Guide

hp OpenView Continuous Access Storage Appliance

Product Version: 5.6.1

Second Edition (September 2003)

Part Number: AA–RTD5B–TE

This guide describes how to configure and manage the HP OpenView Continuous Access Storage Appliance (CASA).

Printed in the U.S.A.

Continuous Access Storage Appliance System Administrator's Guide
Second Edition (September 2003)
Part Number: AA–RTD5B–TE

contents

About this Guide . . . 9
Overview . . . 10
Intended Audience . . . 10
Related Documentation . . . 10
Conventions . . . 11
Document Conventions . . . 11
Text Symbols . . . 11
Equipment Symbols . . . 12
Rack Stability . . . 13
Getting Help . . . 14
HP Technical Support . . . 14
HP Storage Website . . . 14
HP Authorized Reseller . . . 14

1 Overview of CASA . . . 15
About CASA . . . 16
Benefits . . . 16
Supported Operating Systems . . . 16
Installation Options . . . 16
How CASA Works . . . 17
Interfaces . . . 17
Peer Nodes . . . 17
Hosts . . . 17
Storage Virtualization . . . 18
Users and Privileges . . . 18
Local FCP Mirroring . . . 18
Cross-Appliance FCP Mirroring . . . 19
Switch Management . . . 19

2 Optional CASA Features . . . 21
IP Mirroring . . . 22
Point-in-Time Images . . . 22
HP StorageWorks Auto Path . . . 22

3 Using the Interfaces . . . 23
Logging In to the User Interface . . . 24
Setting the Login Window as the Home Page . . . 24

Setting Web Browser Parameters. 25
Navigating the User Interface. 26
User Interface Components 26
About the Home Window 26
Menu Bar Options 28
Using the Command Line Interface 29
Using the CASA Console. 30
About the CASA Console 30
Console Options. 31
Event Viewer 32
Application Logs 32
Security Logs 33
System Logs. 34
Filtering Event Logs. 35
Viewing Event Details 36
Log Viewer. 37
Filtering SANOS Logs. 39

4 Configuring Nodes 41
Adding Display Nodes 42
About Display Nodes. 42
Procedure 42
Deleting Display Nodes 44
Procedure 44
Viewing the Local Node Configuration 45
Adding Peer Nodes. 46
Procedure 46
Deleting Peer Nodes. 47
Procedure 47

5 Configuring Target Ports 49
About Targets 50
Configuring Target Ports 51

6 Managing Hosts 53
About Host Connections. 54
Registering Host Information 55
Command Line Interface 55
Sanreg Utility 55
User Interface. 58
Understanding the HBA Mask Field 59
Deleting Hosts 60

7 Managing LUNs 61
About LUNs 62
Mapping LUNs. 63
About LUN Maps. 63
LUN Map Rules. 63
Creating LUN Maps. 64
Deleting LUN Maps. 67

Resizing LUNs 68
Viewing Resized LUNs 68
Creating LUN Partitions 69
LUN Partition Rules 69
Procedure 70
Deleting LUN Partitions 72
Creating LUN Expansions 73
About LUN Expansion 73
LUN Expansion Rules 74
LUN Expansion Types 74
Procedure 75
Deleting LUN Expansions 77
Viewing LUN Information 78
Navigating the view luns window 78
Viewing LUN History 81
Viewing the LUN Tree 82
Viewing Physical LUNs 82
Viewing Virtual LUNs 85
Adding LUN Location 87
Viewing LUN Statistics 90
Viewing the LUN Matrix 92
Deleting LUNs 94
Rescanning LUNs 95

8 Managing User Access 97
About Roles and Privileges 98
Creating Roles 99
Modifying Roles 101
Modifying User Roles 101
Editing Role Privileges 101
Deleting Roles 103
Creating Users 104
Changing User Passwords 106
Deleting Users 107

9 Managing Local FCP Mirrors 109
About Local FCP Mirrors 110
Local FCP Mirror Details 111
Rules for Source and Target LUNs 111
Creating a Local FCP Mirror 112
About Mirror States 114
Pausing Local FCP Mirrors 116
Resynchronizing Local FCP Mirrors 118
Deleting Local FCP Mirrors 120

10 Managing Cross-Appliance FCP Mirrors 121
About Cross-Appliance FCP Mirrors 122
Creating a Cascaded Configuration 123
Wiring Cascaded Appliances 123
Identifying Cascaded Nodes 124

Exporting LUNs 126
About Exported LUNs 126
Creating Exported LUNs 127
Procedure 127
Viewing Exported LUNs 132
User Interface 132
LUN Matrix Window 132
LUN Details Window 133
Command Line Interface 135
Cross-Appliance Mirror Details 136
Using Exported LUNs for IP Mirrors 136
Creating Cross-Appliance FCP Mirrors 137
About Mirror States 140
Pausing Cross-Appliance FCP Mirrors 141
Resynchronizing Cross-Appliance FCP Mirrors 143
Deleting Cross-Appliance FCP Mirrors 145

11 Managing IP Mirrors 147
About IP Mirrors. 148
Identifying Mirror Relationships 149
About Mirror Relationships 149
About Secondary Paths 149
Dual WAN Configuration 150
Using the Proxy Path 151
Mirror Relationship Example 152
Central WAN Configuration 153
Using the Direct Path. 154
Mirror Relationship Example 155
Adding Mirror Nodes 156
Adding Multiple Mirror Nodes. 158
About Many-to-One IP Mirroring 158
Identifying Mirror Relationships 159
Sources to Target 159
Target to Sources 159
IP Mirror Components 160
Synchronous Mirroring 160
Asynchronous Mirroring 160
About Guaranteed Write Ordering 161
About the Drainer 161
Moving the Queue 161
Creating IP Mirrors. 162
About Mirror States 165
Mirror States for Synchronous Mirrors 166
Synchronous Mirror States and Transitions. 168
Mirrors States for Asynchronous Mirrors 169
Asynchronous Mirror States and Transitions. 172
Adjusting IP Mirror Parameters 173
Pausing IP Mirrors 174
Synchronous Mirrors 174
Asynchronous Mirrors 174

Resynchronizing IP Mirrors 177
Choosing the Resynchronization Direction 178
Identifying the Resynchronization Node 180
Performing a Mirror Resynchronization 180
Deleting IP Mirrors 182
Removing Mirror Nodes 183

12 Managing Point-in-Time Images 185
About Point-in-Time Images 186
About the Backing Store 186
Point-in-Time Image Attributes 186
Ensuring Data Accuracy 188
Creating a Point-in-Time Image Target 189
Creating a Point-in-Time Image 192
Viewing Available Point-in-Time Image Targets 194
Deleting a Point-in-Time Image 195
Deleting a Point-in-Time Image Target 196

13 Using the Switch Interface 197
Registering Switches 198
Managing Multiple Switch Connections 199
Command Types 200
SMS Commands 200
Connect 200
Logout 200
Brocade Switch Commands 202

A Backing Up and Restoring Files 203
Backing Up Files 204
Restoring Files 205

B Using Scripts 207
About Scripts 208
Creating a Script 209
Scripts and Comments 210
Executing a Script 211
Scripts and LAN Connections 211
Using the Perl Module 212

Glossary 213

Index 221

about this guide

This guide provides information to help you manage and configure the HP OpenView Continuous Access Storage Appliance (CASA) in your storage area network (SAN) environment.

“About this Guide” covers the following topics:

- Overview, page 10
- Conventions, page 11
- Rack Stability, page 13
- Getting Help, page 14

Overview

This section covers the following topics:

- Intended Audience
- Related Documentation

Intended Audience

This guide is intended for users familiar with the Continuous Access Storage Appliance and experienced with CASA supported:

- SAN fabric configurations
- Host operating systems
- Storage systems

Related Documentation

Refer to the following documentation for more information about the Continuous Access Storage Appliance:

- *HP OpenView Continuous Access Storage Appliance Release Notes*
- *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide*
- HP OpenView Continuous Access Storage Appliance Online Help

Conventions

Conventions consist of the following:

- Document Conventions
- Text Symbols
- Equipment Symbols

Document Conventions

The document conventions included in Table 1 apply in most cases.

Table 1: Document Conventions

Element	Convention
Cross-reference links	Blue text: Figure 1
Key and field names, menu items, buttons, and dialog box titles	**Bold**
File names, application names, and text emphasis	*Italics*
User input, command and directory names, and system responses (output and messages)	`Monospace font` `COMMAND NAMES` are uppercase monospace font unless they are case sensitive
Variables	*`<monospace, italic font>`*
Website addresses	Blue, underlined sans serif font text: http://www.hp.com

Text Symbols

The following symbols may be found in the text of this guide. They have the following meanings.

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.

Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Equipment Symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings.

Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal safety from electrical shock hazards, do not open this enclosure.

Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal safety from a hot component, allow the surface to cool before touching.

Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal safety from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.

Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal safety or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

Rack stability protects personnel and equipment.

WARNING: To reduce the risk of personal safety or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
- The full weight of the rack rests on the leveling jacks.
- In single rack installations, the stabilizing feet are attached to the rack.
- In multiple rack installations, the racks are coupled.
- Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our website: http://www.hp.com.

HP Technical Support

In North America, call technical support at 1-800-633-3600 (select option 1), available 24 hours a day, 7 days a week.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed questions

HP Storage Website

The HP website has the latest information on this product. Access information about the appliance at

http://h18000.www1.hp.com/storage/software.html.

Scroll to the **Business Continuity** section and choose **HP OpenView Continuous Access Storage Appliance**.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- Elsewhere, see the HP website for locations and telephone numbers: http://www.hp.com.

Overview of CASA

1

This chapter provides an overview of the Continuous Access Storage Appliance and its features. It covers the following topics:

- About CASA, page 16
- How CASA Works, page 17

About CASA

The Continuous Access Storage Appliance enables you to centrally manage storage resources and share and protect critical business information. Authorized HP personnel install the appliance and connect it to your hosts and your storage arrays using Fibre Channel cables.

Benefits

The appliance is designed primarily for heterogeneous data replication and data migration. The benefits of using the appliance include:

- Enterprise-level business continuity capabilities for heterogeneous storage environments
- Cost-effective local and remote data replication
- Increased storage utilization and control of heterogeneous storage capacity
- Online migration between storage arrays
- Increased productivity by enabling more online recovery of data

Supported Operating Systems

The appliance supports all major operating systems (HP-UX, Windows Server 2003, Windows 2000, Windows NT, IBM-AIX, Solaris, and Linux) and distributed network applications, such as Internet access and enterprise resource planning. It is also compatible with HP Command View and HP OpenView software products.

Installation Options

You can choose to have the appliance installed either:

- In the standard HP rack system.
- In the rack of your choice.

Following initial installation and configuration, you can access the appliance using Internet Explorer 5.5 or 6.0 and manage the storage environment from any workstation mapped to the local network. HP also offers support options to assist you with the ongoing service and maintenance of the appliance.

How CASA Works

The appliance has several features that enable you to easily manage your hosts and storage arrays. This section covers the following topics:

- Interfaces
- Peer Nodes
- Hosts
- Storage Virtualization
- Users and Privileges
- Local FCP Mirroring
- Cross-Appliance FCP Mirroring
- Switch Management

See "Optional CASA Features" on page 21 for information on other features that are sold separately.

Interfaces

You can manage the appliance through the user interface or the command line interface (CLI). The user interface includes a multi-node management feature that allows you to manage multiple appliances from a single user interface. You can use the user interface or the CLI to remotely manage the storage network from your desktop or laptop, as long as it is connected to the same network as the appliance. Using the CASA console, you can view information and perform administrative tasks on the appliance nodes.

See "Using the Interfaces" on page 23 for more information.

Peer Nodes

The appliance contains two independent servers, called *nodes*, that act as peers. Both nodes are always active, as opposed to a primary/secondary setup, in which one node is always active and the other node becomes active only when the primary node is down.

If one peer node fails, the other node continues to provide the high-level functionality required by the storage network configuration. This redundancy prevents network downtime if a failure occurs within the appliance.

The target and initiator ports of the appliance nodes support Fibre Channel Protocol. Each port supports a maximum bandwidth of two gigabits per second, the highest FC speed currently available.

See "Configuring Nodes" on page 41 and "Configuring Target Ports" on page 49 for more information.

Hosts

You can use the user interface or the CLI to manage hosts connected to the appliance. Use *sanreg*, an automated utility, to register information about the hosts.

See "Managing Hosts" on page 53 for more information.

Storage Virtualization

The appliance enables hosts to see storage as logical unit numbers (LUNs). A LUN can represent an individual disk or an entire storage array. By creating a virtual storage pool, you can manage and combine LUNs from different storage arrays and make them accessible to any host on the network as virtual disks. Use the appliance to centrally manage all of your storage arrays.

Storage virtualization enables you to:

- Manage storage capacity more easily.
- Execute storage management tasks more efficiently.
- Assign LUN to each host.
- Add hosts or storage arrays to the network without reconfiguring existing storage allocations.

The virtualization features are:

- **LUN partition**—Divide a physical LUN into smaller virtual LUNs and make them available to multiple hosts.
- **LUN expansion**—Combine physical LUNs and present them as one virtual LUN to the hosts.
- **LUN mapping**—Map virtual LUNs to hosts.
- **LUN tree**—Track how LUNs are being used and view the history of each LUN. Use the labeling feature to specify details about your storage arrays.

See "Managing LUNs" on page 61 for more information.

Users and Privileges

You can maintain the security of the appliance by creating custom user roles and assigning them to users as needed.

See "Managing User Access" on page 97 for more information.

Local FCP Mirroring

Local FCP mirroring is data replication between two storage arrays connected to the same appliance through a Fibre Channel Protocol (FCP) link. It enables you to:

- Create local copies of data.
- Improve the speed of data access.
- Facilitate disaster recovery if a system failure occurs.
- Migrate data between storage arrays without interrupting operations.

See "Managing Local FCP Mirrors" on page 109 for more information.

Cross-Appliance FCP Mirroring

Cross-appliance FCP mirroring extends the benefits of local FCP mirroring to storage arrays connected to different appliances. You enable cross-appliance FCP mirroring by using FCP links to connect the appliances in a cascaded configuration.

See "Managing Cross-Appliance FCP Mirrors" on page 121 for more information.

Switch Management

You can use the CLI to manage switches connected to the appliance. Currently, this feature is only supported on Brocade switches.

See "Using the Switch Interface" on page 197 for more information.

Optional CASA Features

2

HP provides optional features to enhance the functionality of your appliance. This chapter covers the following topics:

- IP Mirroring, page 22
- Point-in-Time Images, page 22
- HP StorageWorks Auto Path, page 22

IP Mirroring

IP mirroring is data replication between storage arrays connected to different appliances through an Internet Protocol (IP) link. The unlimited range of the IP link extends the benefits of local and cross-appliance FCP mirroring to appliances separated by great distances. (See "Local FCP Mirroring" and "Cross-Appliance FCP Mirroring" on page 19.) For example, you can use IP mirroring to replicate data between an appliance in New York and an appliance in London.

IP mirroring also supports many-to-one mirroring, which is useful for centralized backup and disaster recovery. You can have up to three appliances mirror data to one appliance.

See "Managing IP Mirrors" on page 147 for more information.

Point-in-Time Images

Point-in-time images are virtual copies of data created at a specific point in time. Creating a point-in-time image does not prevent users from accessing the source data, nor does it affect bandwidth or disk resources. You can view and use a point-in-time image as if it were a physical LUN.

You can use point-in-time images to:

- Create a backup of a LUN while the LUN continues to be updated.
- Create copies of production data for testing purposes.

See "Managing Point-in-Time Images" on page 185 for more information.

HP StorageWorks Auto Path

HP StorageWorks Auto Path is host-based software that provides I/O path failover and load balancing. This software provides a fault-tolerant infrastructure for redundant I/O paths from the hosts to the appliance with automated protection against failures. Auto Path ensures that a failure on one path does not prevent the host from accessing data through another path. Dynamic load balancing improves the overall performance.

Auto Path features include:

- Failure protection between the appliance and the host
- Dynamic load balancing over multiple paths to a storage device, which enables applications to continue running without interruption if a path failure occurs
- Enhanced data availability
- Commands that display device, path, and adapter information.

Refer to the HP StorageWorks Auto Path documentation for more information.

Using the Interfaces

3

You can manage the appliance using the web-based user interface or the command line interface (CLI). You can access either interface from any server connected to the same network as the appliance. The server can be your desktop or laptop or any host attached to the appliance. You can also view information and perform administrative tasks on the appliance nodes using the **hp OV CASA Console** window. You can access the console directly from the appliance.

This chapter covers the following topics:

- Logging In to the User Interface, page 24
- Navigating the User Interface, page 26
- Using the Command Line Interface, page 29
- Using the CASA Console, on page 30
- Event Viewer, page 32
- Log Viewer, page 37

Logging In to the User Interface

You can access the user interface from any server connected to the same network as the appliance. The server can be your desktop or laptop or any host connected to the same network as the appliance.

Note: The user interface supports Internet Explorer version 5.5 or 6.0.

To log in to the user interface:

1. Open the web browser:
   - If set as the default home page, the **login** window for the user interface opens (Figure 1).
   - If the **login** window is not the default home page, enter the appliance name or IP address in the web browser's **Address** field and press **Enter**.

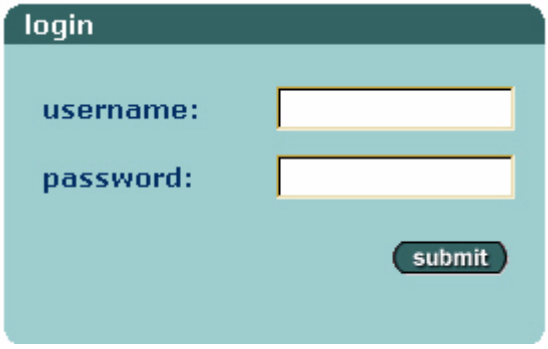

Figure 1: Login window

2. Enter the user name and password and click **submit**.

   The **home** window opens. (See "Navigating the User Interface" on page 29 for more information.)

Setting the Login Window as the Home Page

To set the **login** window as the web browser's home page:

1. Open the web browser.
2. Choose **Tools > Internet Options**.

   The **Internet Options** dialog box opens.
3. Under **Home page**, enter the appliance name or IP address in the **Address** field and click **Use Default**.

Setting Web Browser Parameters

To ensure that the web browser works properly with the appliance's user interface:

1. Open the web browser.
2. Choose **Tools > Internet Options**.

   The **Internet Options** dialog box opens.
3. Click **Settings** in the **Temporary Internet files** field
4. Choose **Every visit to the page** in the **Check for newer versions of stored pages** field.
5. Click **OK**.
6. Close the **Internet Options** dialog box.

Navigating the User Interface

This section describes how to navigate the user interface. It covers the following topics:

- User Interface Components
- About the Home Window
- Menu Bar Options

User Interface Components

Each user interface window has a common format and structure (Figure 2).

- Use the menu bar (**A**) to navigate to any option.
- When you choose an option from the menu bar, choose submenu options from the mouseover list directly under the menu bar (**B**) or from the list on the left side of the window (**C**).
- The selected option is displayed in the main area of the window (**D**).
- The multi-node management feature (**E**) allows you to manage multiple appliances from a single user interface. When you choose a node from the drop-down list, you can view and manage that node from the current user interface window. You do not have to log out and log in on the other node. (See "Adding Display Nodes" on page 42 for more information.)
- To access online help for the user interface, click **HELP** (**G**).
- To exit the user interface, click **LOGOUT** (**G**).

About the Home Window

After you log in to the user interface, the **home** window opens (Figure 2) and displays the following:

- Name and IP address of the appliance node you are currently using (**D1**)
- Name and IP address of the peer node (the other node within the appliance)
- Names and IP addresses of any mirror nodes
- Data that the appliance is managing (**D2**), including:
  - LUNs
  - Hosts
  - LUN maps
  - Partitions
  - Expansions
  - Mirrors
- Real-time information about appliance node activity (**F**)

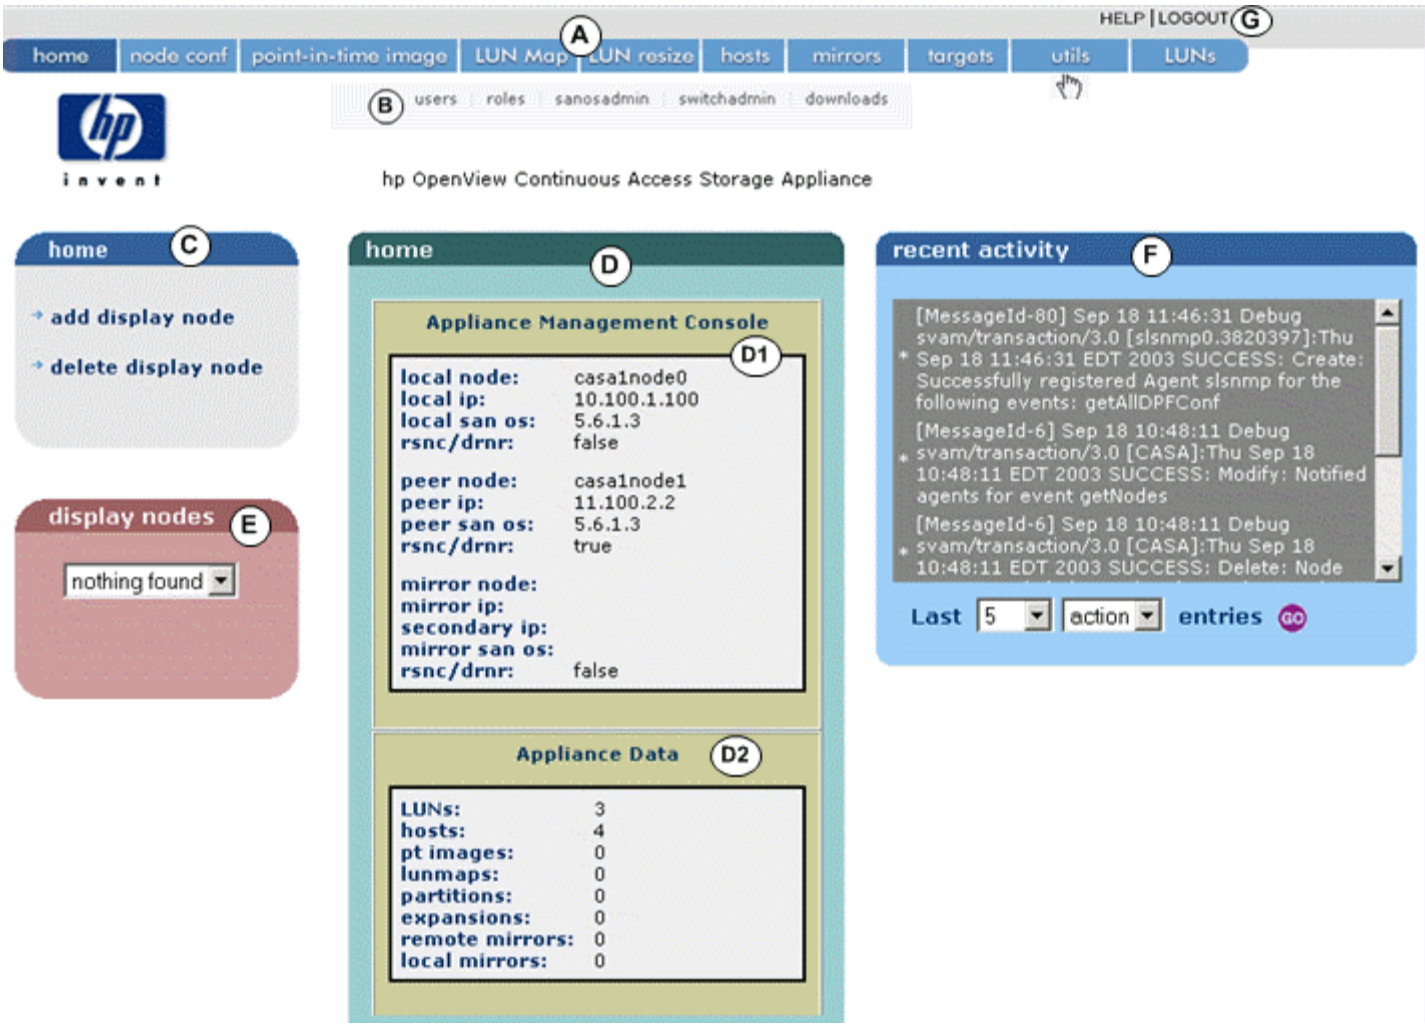

Figure 2: Home window

Menu Bar Options

Table 2 lists the menu bar options on the **home** window and identifies where to find information about them.

Table 2: Menu bar options

Menu Bar Option	For More Information
home	See "About the Home Window" on page 27.
node conf	See "Configuring Nodes."
point-in-time image	See "Managing Point-in-Time Images."
LUN Map	See "Managing LUNs."
LUN resize	See "Managing LUNs."
hosts	See "Managing Hosts."
mirrors	See the appropriate mirroring chapter: ■ "Managing Local FCP Mirrors" ■ "Managing Cross-Appliance FCP Mirrors" ■ "Managing IP Mirrors"
targets	See "Configuring Target Ports."
utils	See "Managing User Access."
LUNs	See "Managing LUNs."

Using the Command Line Interface

You can also manage the appliance using the CLI. You can access the CLI commands from the CASA console or by choosing **utils > sanosadmin** from the user interface.

Refer to the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide* for detailed information about the CLI. The reference guide provides an overview of the CLI, including command types and access methods, and lists the syntax for each CLI command.

For most functions, you can use either the CLI or the user interface. However, there are some functions that are only accessible from the CLI and some functions that are only accessible from the user interface. This is noted, where applicable, in this guide and in the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide*.

Using the CASA Console

This section explains how to use the **hp OV CASA Console** window. It covers the following topics:

- About the CASA Console
- Console Options

About the CASA Console

The **hp OV CASA Console** window opens when you access the appliance's monitor. Use the console to perform administrative tasks and view information about the node you are currently accessing (Figure 3).

You can view the following information for the appliance node:

- **CPU/MEM/NET**—The central processing unit, memory, and network meters are updated dynamically.
- **system info**—Hardware information, including the appliance name, available memory and disk space, uptime, and appliance node IP address.
- **CASA info**—Software information, including the appliance software (called SANOS), management software, and HTTP server software versions used.
- **Current date and time—**Mouseover the time to display the date.

Figure 3: hp OV CASA Console window

Console Options

The console options are:

- **command view**—Access the Secure Path application. Refer to your Secure Path documentation for more information.
- **sanos admin**—Choose **initialize** to create an administrative user. Choose **launch** to access the CLI.
- **event viewer**—Access application, security, and system logs. (See "Event Viewer" for more information.)
- **log viewer**—Access the appliance software (SANOS) logs. (See "Log Viewer" for more information.)
- **backup\restore**—Access scripts to back up and restore the appliance configuration data. (See "Backing Up and Restoring Files" on page 203 for more information.)
- **shutdown**—Shut down the node.

Event Viewer

Use the event viewer to application, security, and system logs. This sections covers the following topics:

- Application Logs
- Security Logs
- System Logs
- Filtering Event Logs
- Viewing Event Details

Application Logs

The application log contains software events logged by the appliance's operating system.

To view the application log, choose **event viewer** from the **hp OV CASA Console** window. The **application log** window opens (Figure 4).

event viewer - application log

Log View

Type	Date	Time	Source	Category	Event	User	CASA
Information	12/18/2002	9:44:46	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/18/2002	7:38:47	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/17/2002	11:41:25	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/17/2002	11:07:03	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/17/2002	9:02:11	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/17/2002	8:58:20	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/17/2002	8:42:43	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/17/2002	7:52:45	CBRegCap	None	0	N/A	MINOTTJOYCE
Error	12/17/2002	7:52:44	LPR Print Monitor	None	2007	N/A	MINOTTJOYCE
Information	12/16/2002	14:23:17	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/16/2002	8:33:09	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/16/2002	7:57:54	CBRegCap	None	0	N/A	MINOTTJOYCE
Error	12/16/2002	7:57:53	LPR Print Monitor	None	2007	N/A	MINOTTJOYCE
Information	12/13/2002	7:49:05	CBRegCap	None	0	N/A	MINOTTJOYCE
Error	12/13/2002	7:49:04	LPR Print Monitor	None	2007	N/A	MINOTTJOYCE
Information	12/12/2002	8:35:06	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/12/2002	8:32:42	MsiInstaller	None	11707	N/A	MINOTTJOYCE
Information	12/12/2002	8:08:15	MsiInstaller	None	11724	N/A	MINOTTJOYCE
Information	12/12/2002	7:56:03	CBRegCap	None	0	N/A	MINOTTJOYCE
Information	12/11/2002	10:04:06	CBRegCap	None	0	N/A	MINOTTJOYCE
Error	12/11/2002	10:04:05	LPR Print Monitor	None	2007	N/A	MINOTTJOYCE
Information	12/11/2002	9:27:15	CBRegCap	None	0	N/A	MINOTTJOYCE

Figure 4: Application log window

Security Logs

The security log records all login attempts and resource-related events, such as creating, opening, or deleting files. Event types success audit and failure audit. A success audit identifies security access attempts that succeeded. A failure audit identifies security access attempts that failed.

To view the security log:

1. Choose **event viewer** from the **hp OV CASA Console** window.

   The **application log** window opens (Figure 4).

2. From the event viewer task bar, choose **Log > Security**.

   The **security log** window opens (Figure 5).

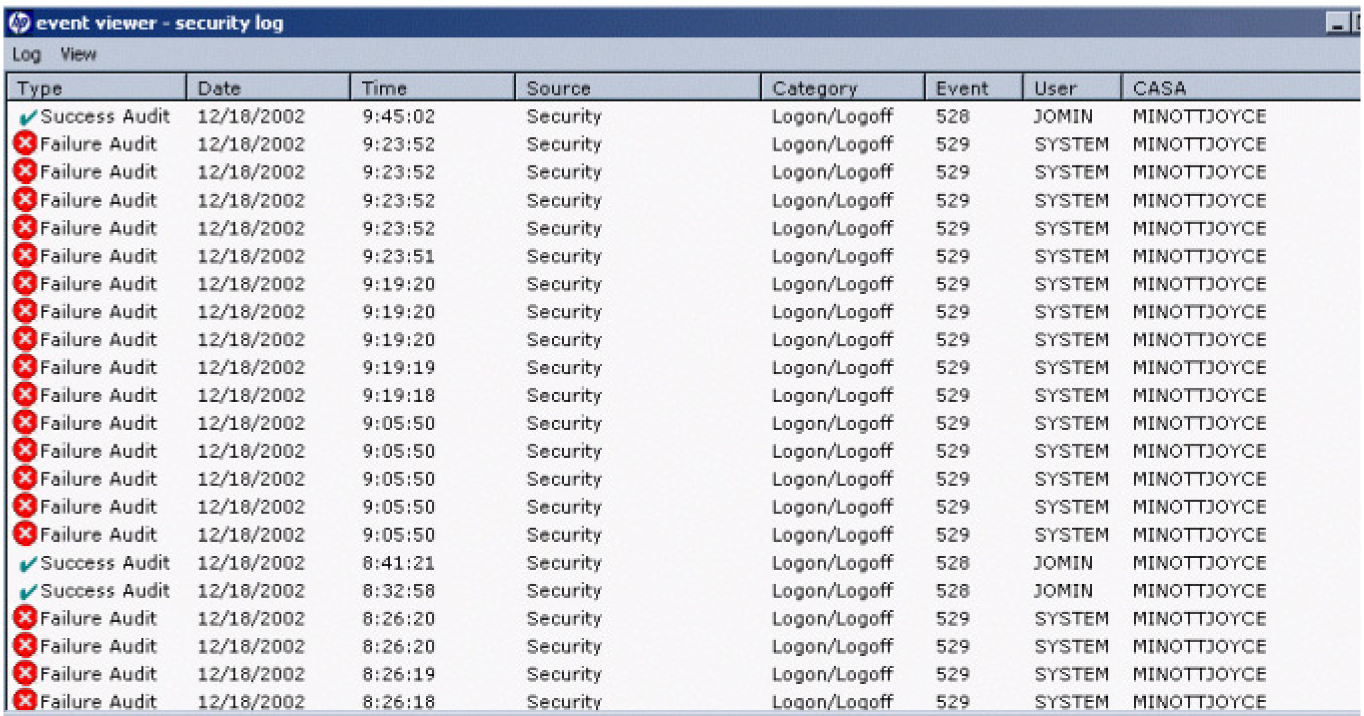
event viewer - security log
Log View

Type	Date	Time	Source	Category	Event	User	CASA
Success Audit	12/18/2002	9:45:02	Security	Logon/Logoff	528	JOMIN	MINOTTJOYCE
Failure Audit	12/18/2002	9:23:52	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:23:52	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:23:52	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:23:52	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:23:51	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:19:20	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:19:20	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:19:20	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:19:19	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:19:18	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:05:50	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:05:50	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:05:50	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:05:50	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	9:05:50	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Success Audit	12/18/2002	8:41:21	Security	Logon/Logoff	528	JOMIN	MINOTTJOYCE
Success Audit	12/18/2002	8:32:58	Security	Logon/Logoff	528	JOMIN	MINOTTJOYCE
Failure Audit	12/18/2002	8:26:20	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	8:26:20	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	8:26:19	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE
Failure Audit	12/18/2002	8:26:18	Security	Logon/Logoff	529	SYSTEM	MINOTTJOYCE

Figure 5: Security log window

System Logs

The system log contains hardware events logged by the appliance's operating system.

To view the system log:

1. Choose **event viewer** from the **hp OV CASA Console** window.
   The **application log** window opens (Figure 4).
2. From the event viewer task bar, choose **Log > System**.
   The **system log** window opens (Figure 6).

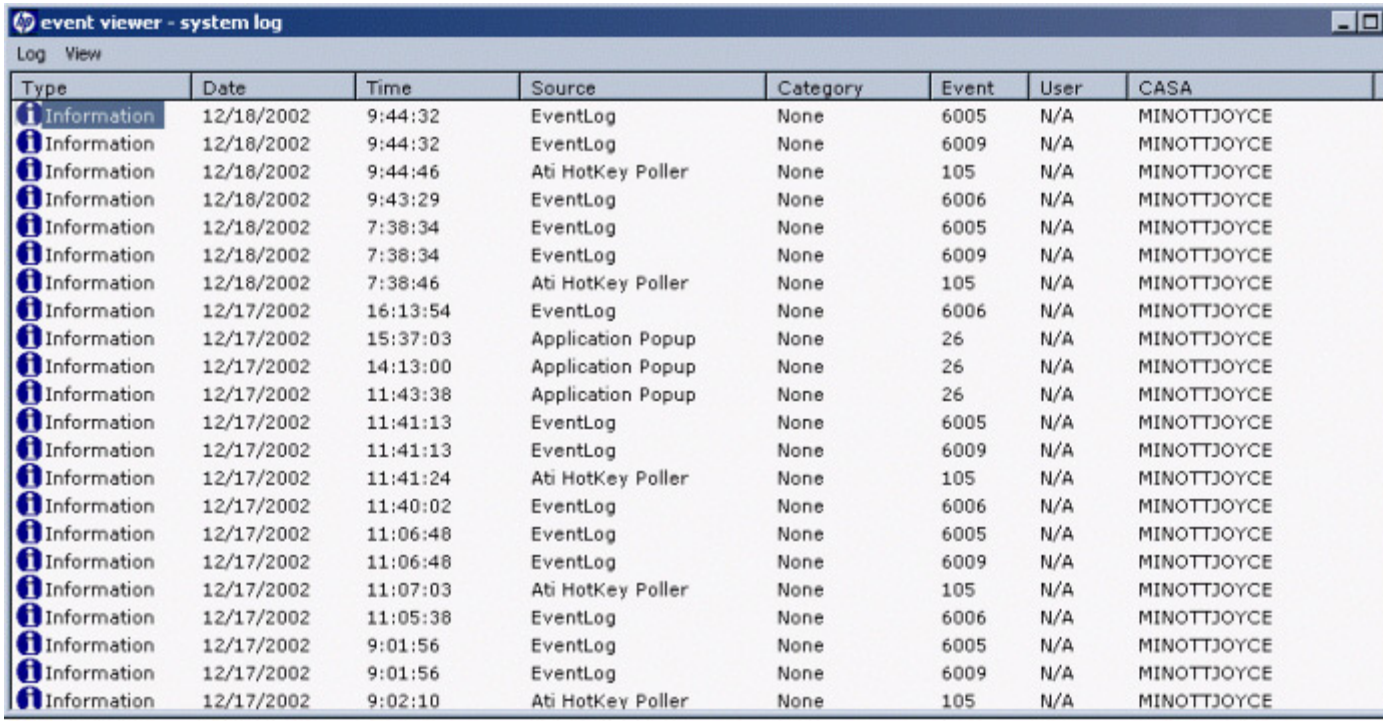

Type	Date	Time	Source	Category	Event	User	CASA
Information	12/18/2002	9:44:32	EventLog	None	6005	N/A	MINOTTJOYCE
Information	12/18/2002	9:44:32	EventLog	None	6009	N/A	MINOTTJOYCE
Information	12/18/2002	9:44:46	Ati HotKey Poller	None	105	N/A	MINOTTJOYCE
Information	12/18/2002	9:43:29	EventLog	None	6006	N/A	MINOTTJOYCE
Information	12/18/2002	7:38:34	EventLog	None	6005	N/A	MINOTTJOYCE
Information	12/18/2002	7:38:34	EventLog	None	6009	N/A	MINOTTJOYCE
Information	12/18/2002	7:38:46	Ati HotKey Poller	None	105	N/A	MINOTTJOYCE
Information	12/17/2002	16:13:54	EventLog	None	6006	N/A	MINOTTJOYCE
Information	12/17/2002	15:37:03	Application Popup	None	26	N/A	MINOTTJOYCE
Information	12/17/2002	14:13:00	Application Popup	None	26	N/A	MINOTTJOYCE
Information	12/17/2002	11:43:38	Application Popup	None	26	N/A	MINOTTJOYCE
Information	12/17/2002	11:41:13	EventLog	None	6005	N/A	MINOTTJOYCE
Information	12/17/2002	11:41:13	EventLog	None	6009	N/A	MINOTTJOYCE
Information	12/17/2002	11:41:24	Ati HotKey Poller	None	105	N/A	MINOTTJOYCE
Information	12/17/2002	11:40:02	EventLog	None	6006	N/A	MINOTTJOYCE
Information	12/17/2002	11:06:48	EventLog	None	6005	N/A	MINOTTJOYCE
Information	12/17/2002	11:06:48	EventLog	None	6009	N/A	MINOTTJOYCE
Information	12/17/2002	11:07:03	Ati HotKey Poller	None	105	N/A	MINOTTJOYCE
Information	12/17/2002	11:05:38	EventLog	None	6006	N/A	MINOTTJOYCE
Information	12/17/2002	9:01:56	EventLog	None	6005	N/A	MINOTTJOYCE
Information	12/17/2002	9:01:56	EventLog	None	6009	N/A	MINOTTJOYCE
Information	12/17/2002	9:02:10	Ati HotKey Poller	None	105	N/A	MINOTTJOYCE

Figure 6: System log window

Filtering Event Logs

Use the **event log filter** window to specify the logs you want to view. From any log window, choose **View > Filter Events**. The **event log filter** window opens (Figure 7). You can filter event logs by:

- Date and time
- Event type (Information, Warning, Error, Success Audit, Failure Audit)
- Source
- Category
- Event ID
- User
- Computer

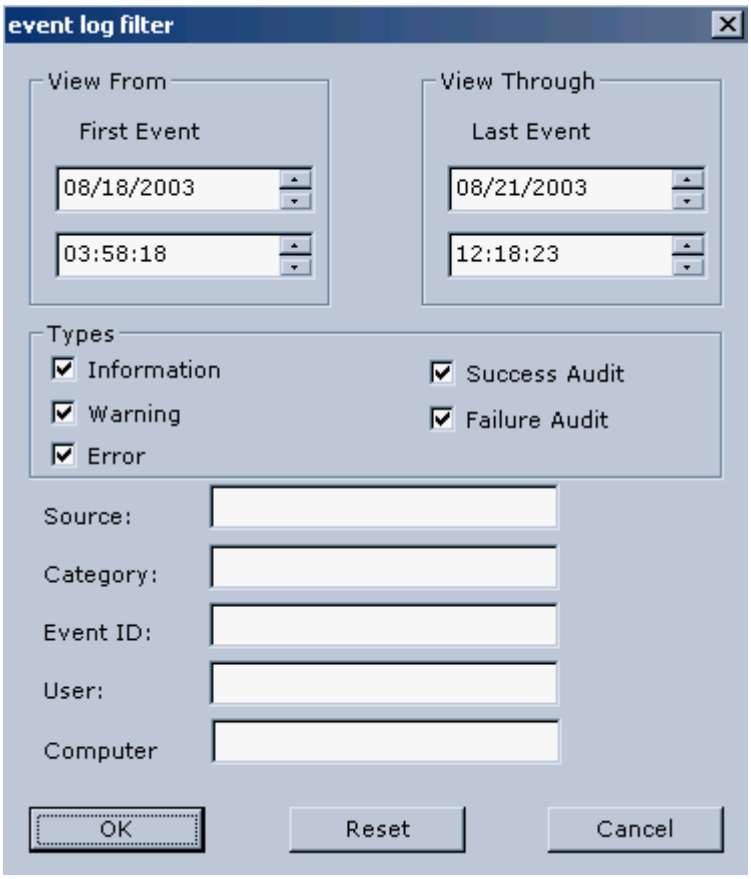

Figure 7: Event log filter window

Viewing Event Details

From any log window, double-click an event to view more detail. The **event detail** window opens (Figure 8).

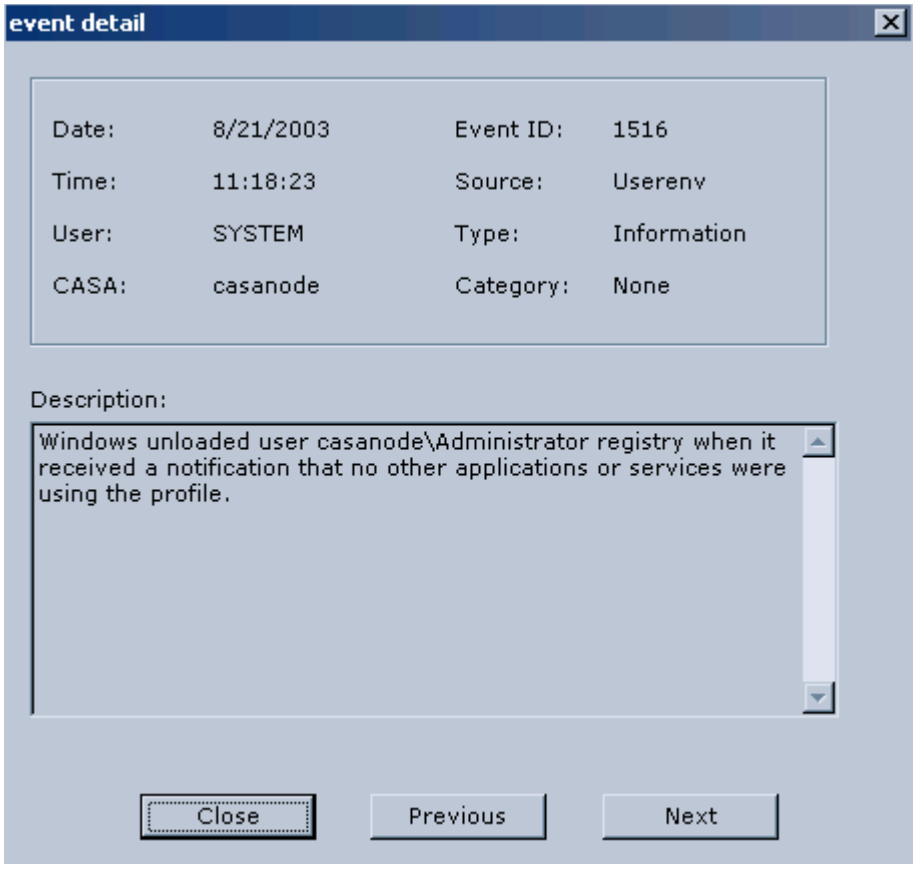

Figure 8: Event detail window

Log Viewer

Use the log viewer to view logs generated by the appliance software (called SANOS).

To view logs of all appliance software activity, choose **log viewer > sanos** from the **hp OV CASA Console** window. The **sanos log** window opens (Figure 9).

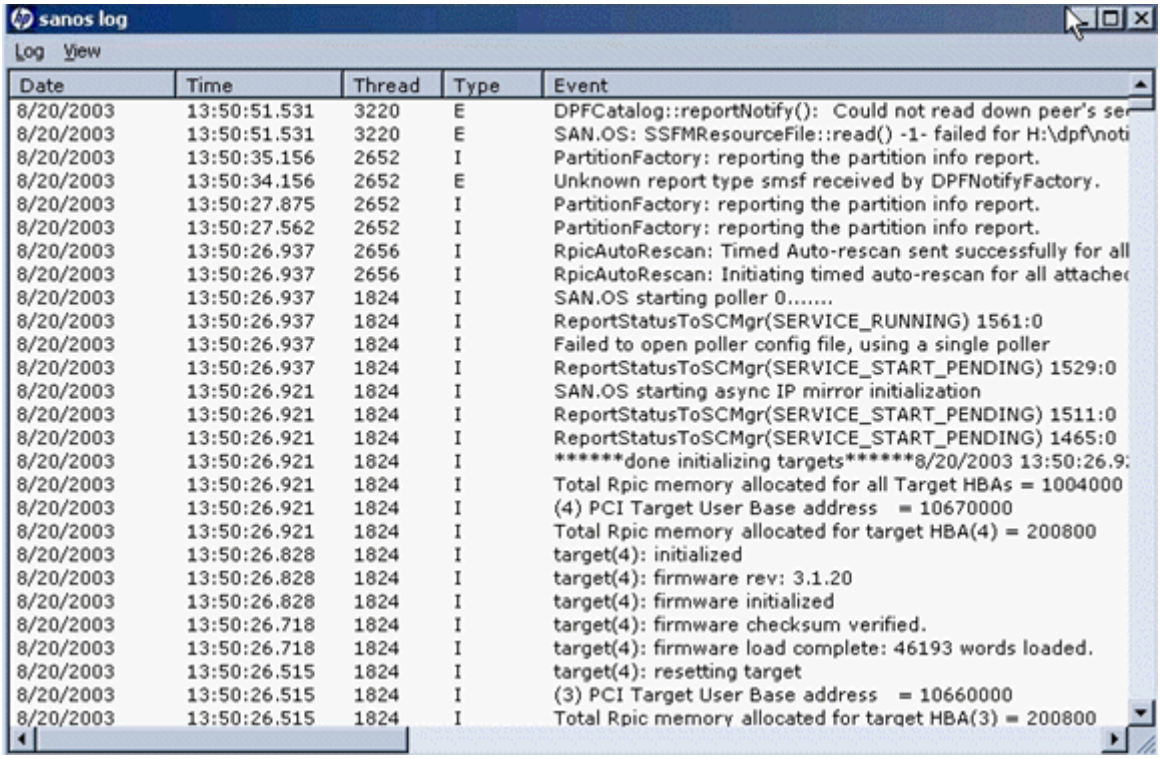

Figure 9: sanos log window

To view logs of all appliance software management activity, choose **log viewer > sanos management** from the **hp OV CASA Console** window. The **sanos management** window opens (Figure 10).

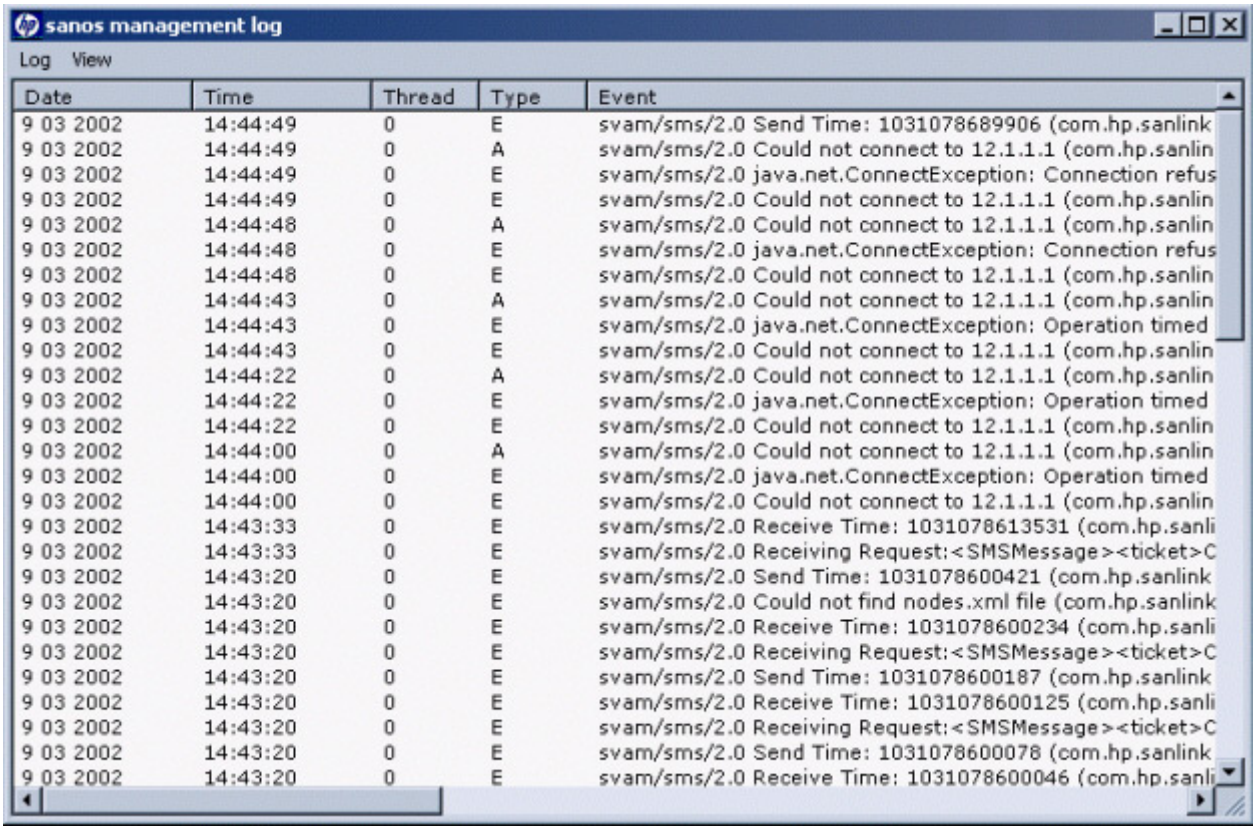

Figure 10: sanos management log

Filtering SANOS Logs

Use the **log filter** window to specify the appliance software log information you want to view (Figure 11). You can filter SANOS logs by:

- Date and time
- Type (Info, Error, Alert, T1, T2, T3)
- Thread

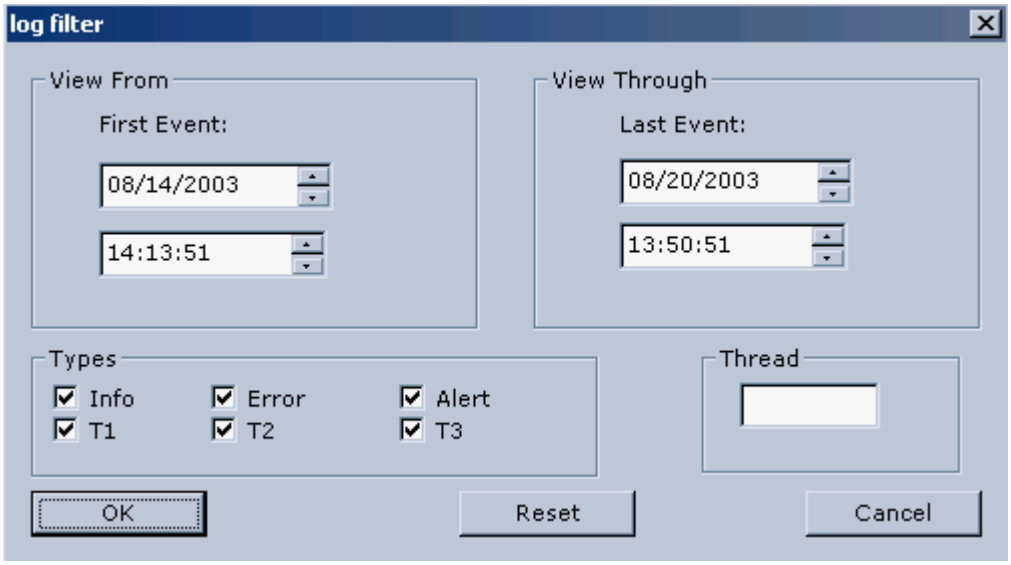

Figure 11: Log filter window

Configuring Nodes

4

This chapter describes how to configure and manage the appliance nodes. It covers the following topics:

- Adding Display Nodes, page 42
- Deleting Display Nodes, page 44
- Viewing the Local Node Configuration, page 45
- Adding Peer Nodes, page 46
- Deleting Peer Nodes, page 47

Adding Display Nodes

To use the multi-node management feature of the user interface, you must add each appliance nodes as a display node.

About Display Nodes

To add a display node, you must specify:

- **Node ip/name**—The public network IP address (or name) of the display node you want to add. Each appliance node has three network addresses: the local LAN address, the peer node address, and the public network address. The public network address identifies the node to other appliances on the network.

Note: HP recommends that you use the node's IP address instead of the node name. If you use the node name, you must ensure that each network connection (LAN, peer, and public) has a unique name. If all three network connections are identified by the same name, you cannot ensure that you are using the public network connection.

- **Add current node to new node**—The *current node* is the node you are currently using. The *new node* is the node you are adding as a display node. If you choose this option, the current node becomes a display node of the new node.

  For example, you are using node A and you want to add node B as a display node. When you select the **Add current node to new node** option, you add node A as a display node from node B.

Procedure

Note: You can only add a display node from the user interface.

To add a display node:

1. Choose **home > add display node**.

   The **create display node** window opens (Figure 12).
2. Enter the display node's IP address or name in the **node ip/name** field.
3. Select the **add current node to new node** check box, if applicable.
4. Click **submit**.

   The **view display nodes** window opens (Figure 13).
5. Repeat steps 1 through 4 to add another display node.

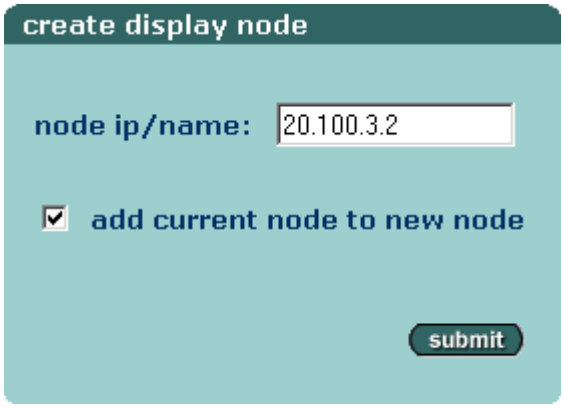

Figure 12: Create display node window

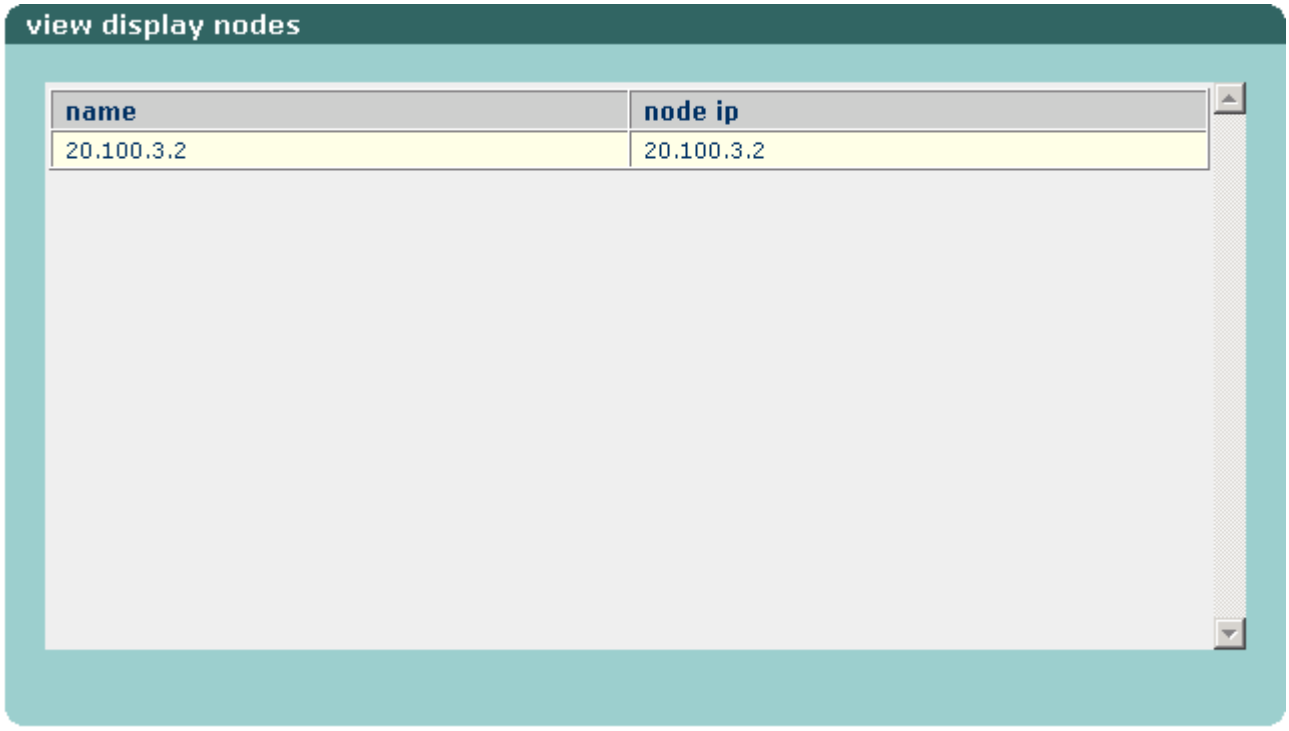

Figure 13: View display nodes window

Deleting Display Nodes

Deleting a display node does not delete it from the network. To access a display node that you deleted, you must open a new user interface window and log in to that node.

Procedure

To delete a display node:

1. Choose **home > delete display node**.

   The **delete display nodes** window opens (Figure 14).

2. Choose the display node you want to delete and click **submit**.

   The **view display nodes** window opens. The display node you deleted is not included.

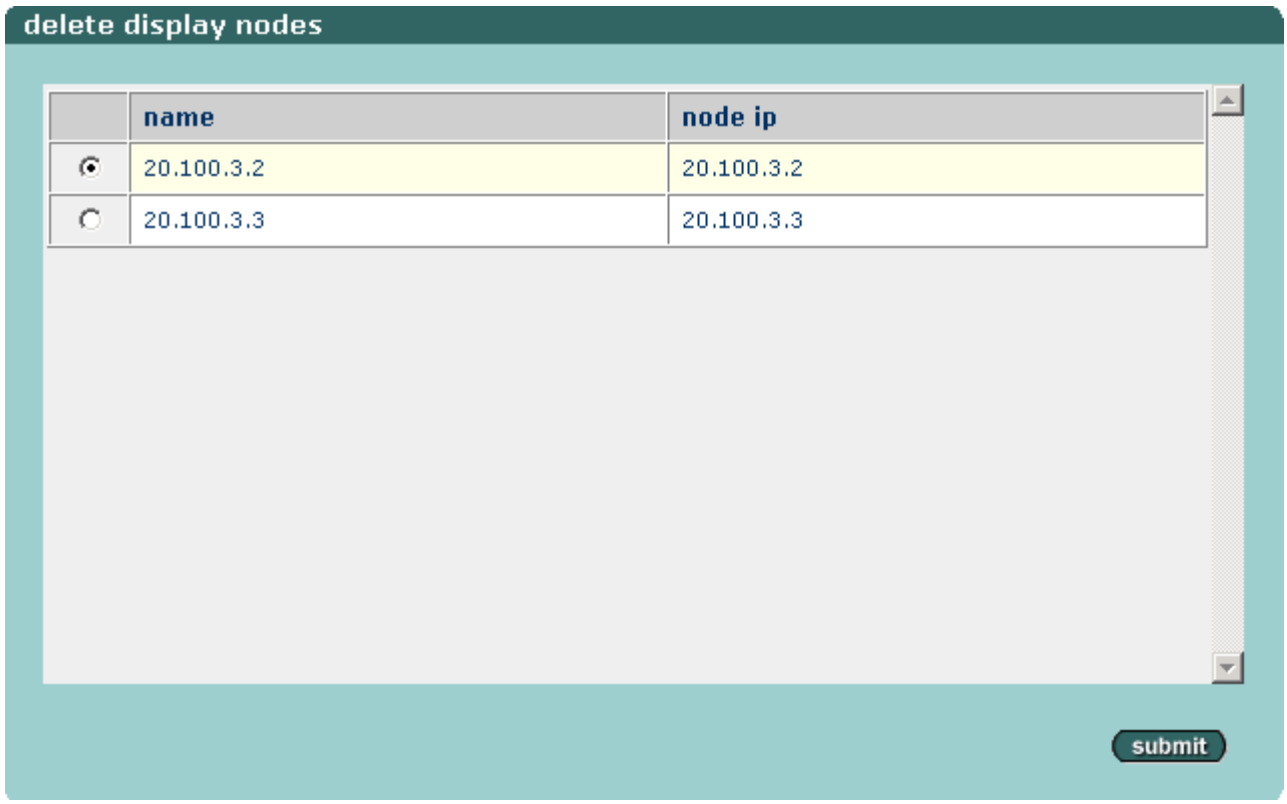

Figure 14: Delete display nodes window

Viewing the Local Node Configuration

To view the local node's configuration, such as its name, IP address and appliance software version, choose **node conf** from the menu bar. The **appliance configuration** window opens (Figure 15). This window also shows the names and IP addresses of the peer node and mirror node (if applicable).

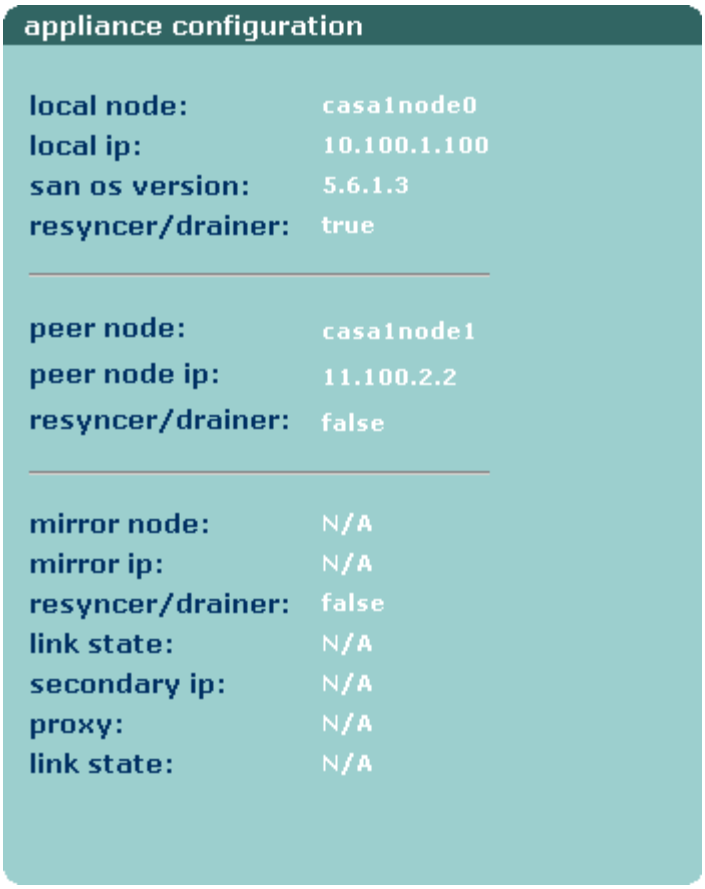

Figure 15: Appliance configuration window

Adding Peer Nodes

HP service personnel add peer nodes when they install the appliance. You will have to re-add peer nodes only if you replace a failed node.

Adding a peer node establishes the connection between the two nodes in an appliance. This connection enables the appliance nodes to communicate. There is only one peer node relationship per appliance node. Node 0 is the peer node of node 1; node 1 is the peer node of node 0.

Each appliance node has three network addresses: the local LAN address, the peer node address, and the public network address. The peer node address identifies the node to its peer.

Note: HP recommends that you use the node's IP address instead of the node name. If you use the node name, you must ensure that each network connection (LAN, peer, and public) has a unique name. If all three network connections are identified by the same name, you cannot ensure that you are using the peer node address.

If you want to change the name or IP address of an appliance node, contact HP support for assistance.

Procedure

To add a peer node:

1. Choose **node conf > add peer node**.

   The **add peer node** window opens (Figure 16).
2. Enter the peer node's name or IP address in the **peer name/IP** field.
3. Click **submit**.

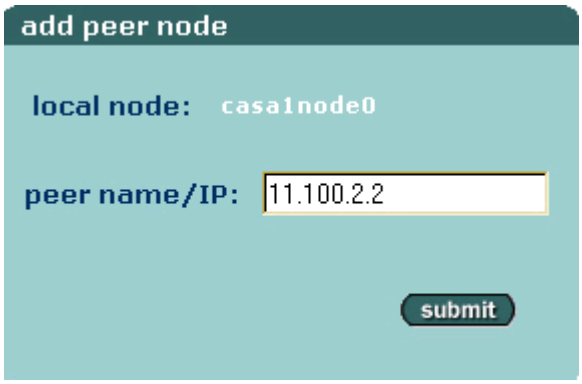

Figure 16: Add peer node window

Deleting Peer Nodes

If you must replace a failed node, first delete the peer node from the appliance's configuration.

You cannot delete a peer node if you have define local and remote mirror nodes or while mirroring data. You must stop mirroring and delete the mirror nodes before deleting the peer node from the configuration. If you do not complete these steps, the user interface displays the following error message:

```
Peer node cannot be removed while local or remote mirrors are
defined. Peer node cannot be removed while mirroring.
```

Refer to "Removing Mirror Nodes" on page 183 and "Deleting IP Mirrors" on page 182 for more information.

Procedure

To delete a peer node:

1. Choose **node conf > remove peer node**.

   The **remove peer node** window opens, displaying the peer node's information (Figure 17).

2. Click **delete**.

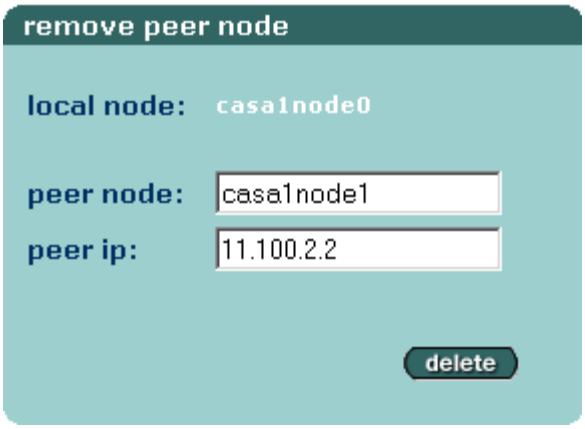

Figure 17: Remove peer node window

Configuring Target Ports

5

This chapter describes how to configure the target ports for each appliance node. The target ports communicate with the host bus adapters (HBAs) of each host that is connected to the appliance. HP service personnel will configure the target port settings for your network when they install the appliance. Unless you make a change to your network, the target port settings should not change.

Caution: Do not change target port settings without assistance from HP service or support personnel. An incorrect setting can cause the appliance to fail.

This chapter covers the following topics:

- About Targets, page 50
- Configuring Target Ports, page 51

About Targets

Each appliance node has six target ports. Using Fibre Channel cables, the target ports and the HBAs of a host are either:

- Directly connected
- Connected through a fabric switch

The Fibre Channel mode (FC mode) determines how the target ports and the HBAs communicate through the Fibre Channel cable connections. The FC mode options are:

- **Loop**—Allows bi-directional communication between ports within a circuit. Use **Loop** when the hosts are directly connected to the appliance.
- **Point to Point** (default)—The target port and the HBA are directly connected through a fabric, enabling them to be connected through a single link. Use **Point to Point** when the hosts and the appliance nodes are connected through a switch.
- **Loop Preferred**—Target ports search for **Loop** mode first and **Point to Point** mode second.
- **Point to Point Preferred**—Target ports search for **Point to Point** mode first and **Loop** mode second.

Configuring Target Ports

To configure a target port:

1. Choose **targets** from the menu bar.

   The **configure targets** window opens (Figure 18).
2. Select the target you want to configure.
3. Select one of the following options in the **fc mode** field:
   - **Loop**
   - **Point to Point**
   - **Loop Preferred**
   - **Point to Point Preferred**
4. Click **set/reset** to save and initialize your selection.
5. To return to the default target settings, click **reset**.

Note: The **node name** and **port name** fields are automatically populated with the node's worldwide name and target port's worldwide name, respectively.

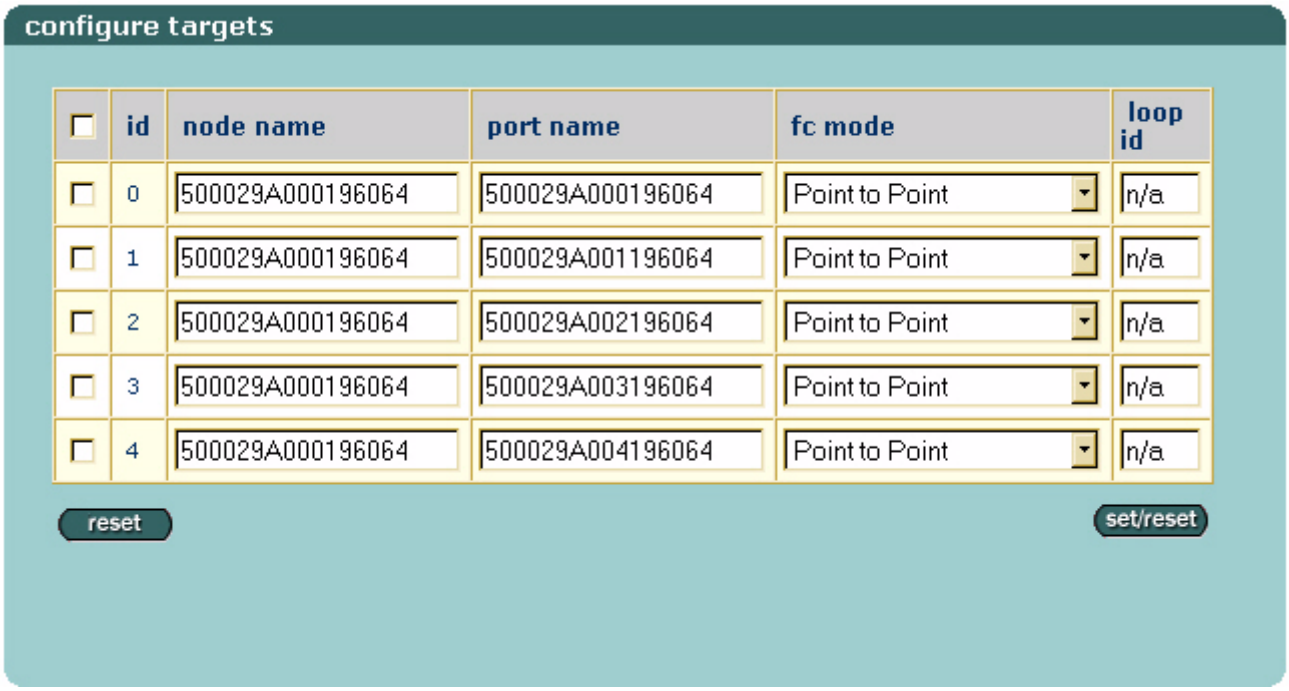

Figure 18: Configure targets window

Managing Hosts

6

The appliance can connect to a variety of hosts. HP service personnel establish the connection between your hosts and the appliance when they install the appliance. Once your hosts are connected to the appliance, you can use the appliance to manage how the hosts access information.

Caution: Only HP service personnel should connect additional hosts or remove faulty hosts.

This chapter covers the following topics:

- About Host Connections, page 54
- Registering Host Information, page 55
- Deleting Hosts, page 60

About Host Connections

When HP service personnel first connect your hosts to the appliance, the appliance software uses specific information to identify each host. In the user interface, choose **hosts** from the menu bar to view this information on the **view hosts** window (Figure 19):

- **id**—The appliance software assigns a number to each host, identifying the order in which the appliance detects hosts.
- **host name**—This field displays the Fibre Channel worldwide name when the host is first connected to the appliance. Once the host is registered, this field displays the name you assigned to the host in its network properties (such as ABC_host_1).
- **wwn**—The Fibre Channel worldwide name of the host.
- **os**—This field displays **UNKNOWN** when the host is first connected to the appliance. Once the host is registered, this field displays the operating system running on the host.
- **hba mask**—The appliance target port to which each host HBA is connected. The appliance software uses a binary calculation to display this number. See "Understanding the HBA Mask Field" on page 59 for more information.
- **node id**—The appliance node (0 or 1) to which the host is connected. Typically, each host has a minimum of two host bus adapters, so that the host has two paths to the appliance, enabling failover.

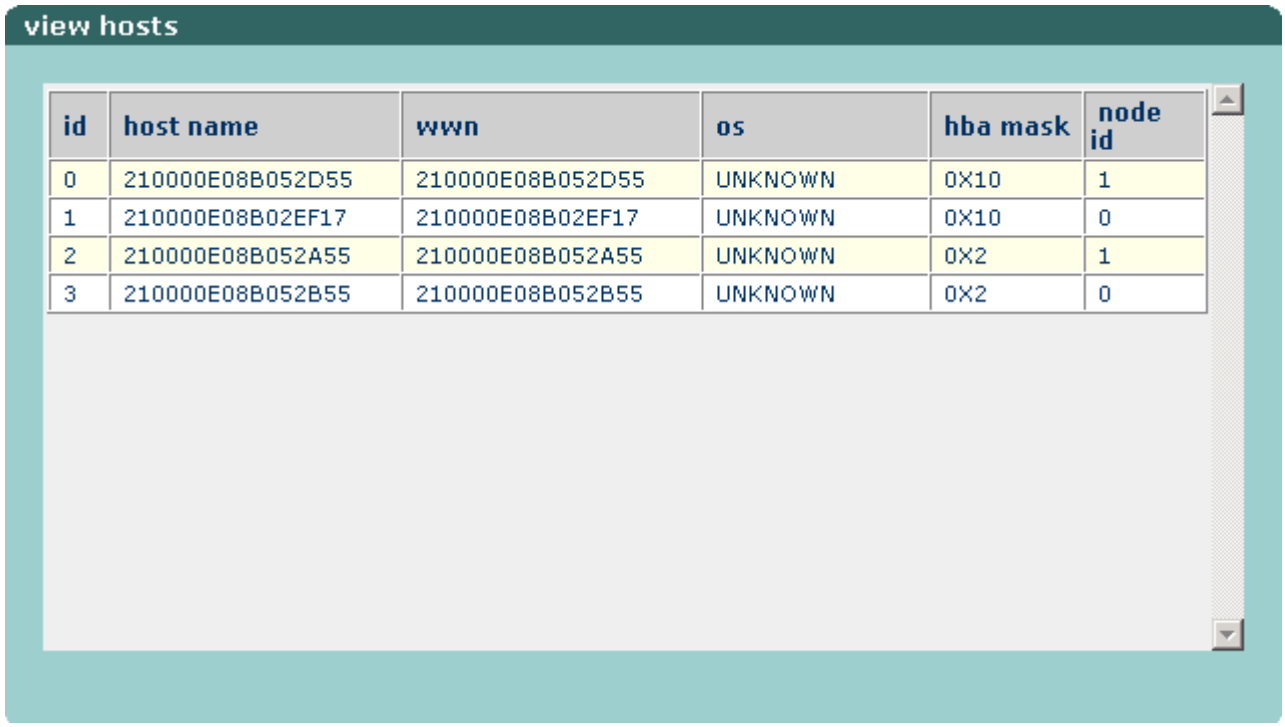

Figure 19: View hosts window

Registering Host Information

You can register host information using:

- Command Line Interface
- Sanreg Utility
- User Interface

Command Line Interface

Refer to the "Host Commands" chapter in the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide* for the appropriate commands.

Sanreg Utility

The `sanreg` utility automatically registers each host. The utility detects and then updates the host's information that displays in the **view hosts** window. HP recommends that you use the `sanreg` utility to register host information.

Note: You must run this utility before you assign storage to a host.

To run the `sanreg` utility:

1. Access the appliance's user interface from the host.
2. Choose **utils > downloads**.

   The **downloads** window opens (Figure 20).
3. Choose **sanreg** for the host's operating system in the **select file to download** field.
4. Click **downloads**.

   When prompted, save the **sanreg** utility to your host's hard disk. Do not download the utility directly from the **downloads** window.
5. Open a command window. Set the current directory to the directory in which you saved `sanreg`, then enter the following command:

   ```
   sanreg
   ```

   The utility displays a message that the host registration was successful (Figure 21).
6. Close the command window and access the appliance's user interface from the host.
7. Choose **hosts** from the menu bar.

   The **view hosts** window opens (Figure 22).

   If the host name includes a number in parentheses, the number indicates the specific host HBA through which the host is connected to the appliance. The number of HBAs includes the adapters that are built into the host. This number ensures that you are mapping LUNs to all adapters in the host.
8. Repeat steps 1 through 7 for each host.

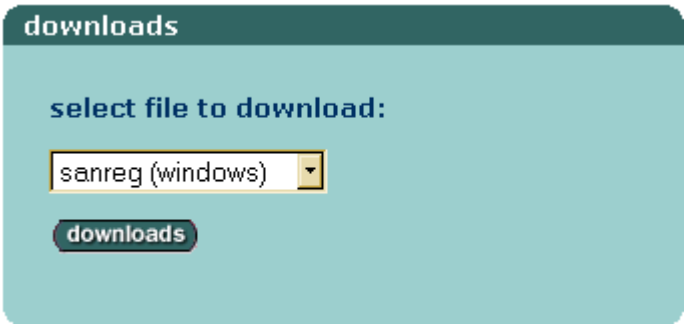

Figure 20: Downloads window

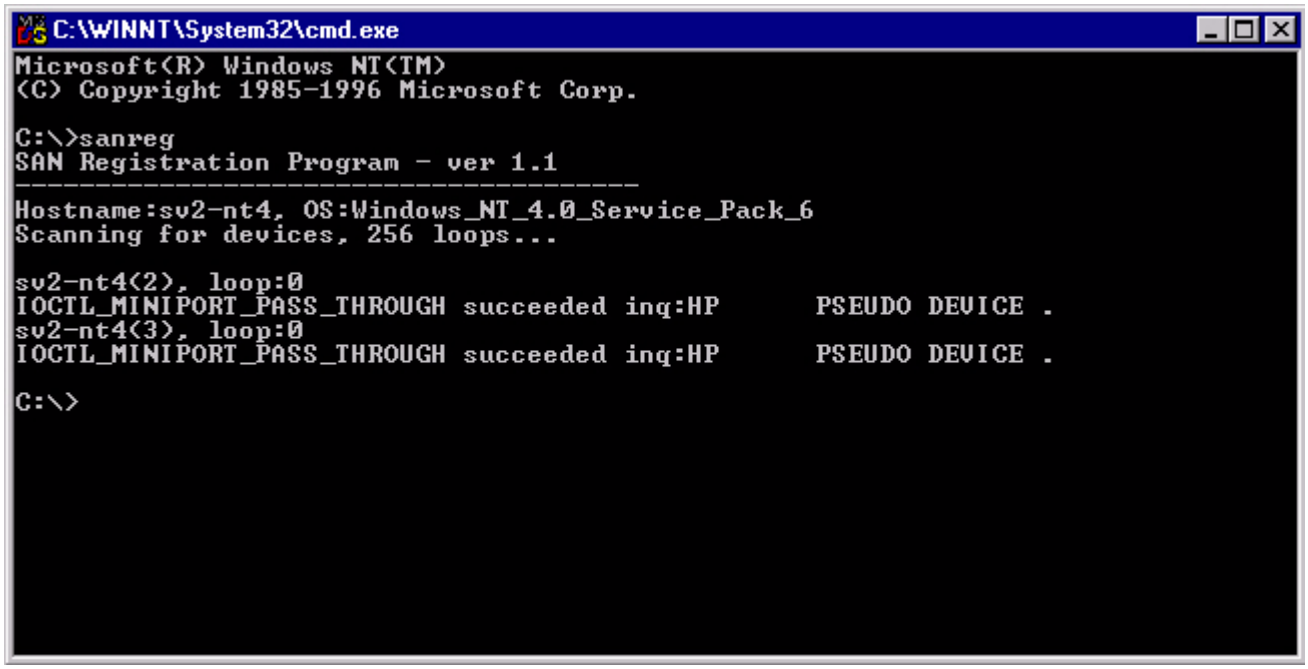

Figure 21: Successful registration message

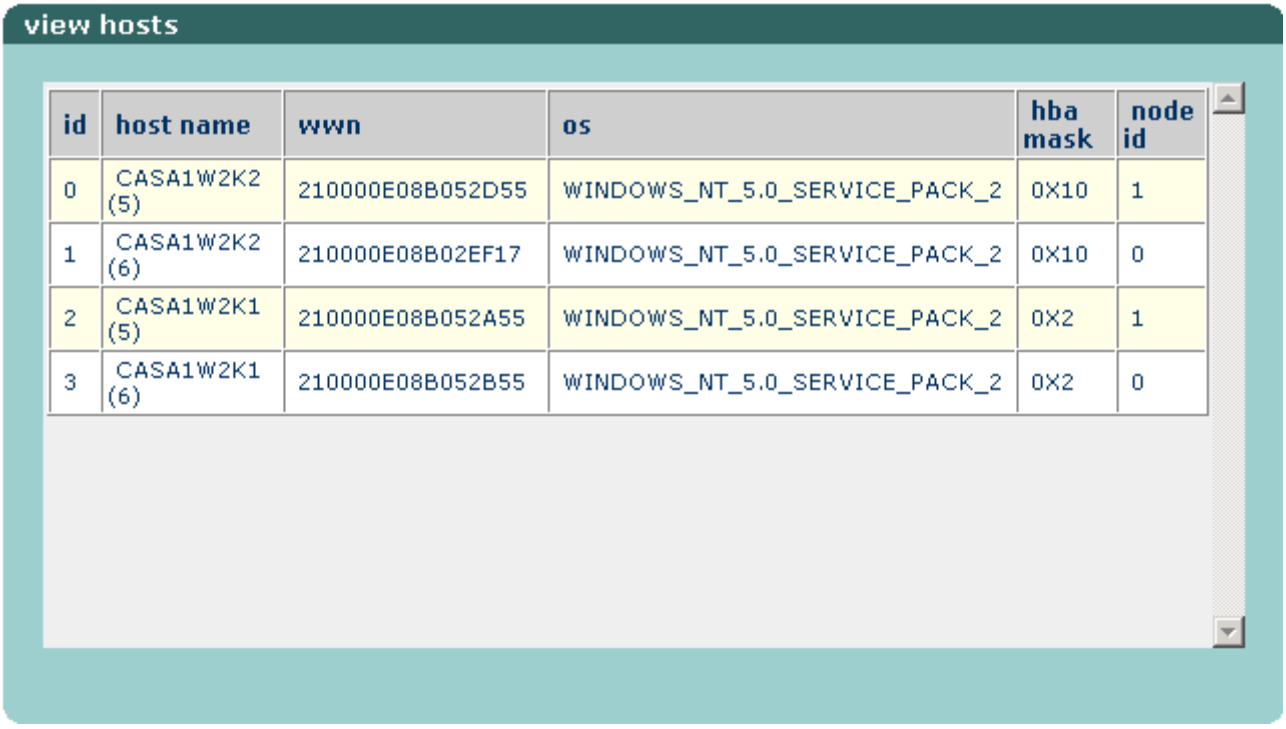

Figure 22: View hosts window (after running the sanreg utility)

User Interface

You can manually register host information using the user interface. HP recommends that you use the `sanreg` utility because it automatically detects and updates the information, eliminating potential errors. (See "Sanreg Utility" on page 55.)

If you use the user interface, you must know exactly how the host HBAs and appliance target ports are connected, particularly if you are using fabric switches and zones. If you select the incorrect target path, you cannot map data properly. If you enter the Fibre Channel worldwide name incorrectly, the appliance cannot communicate with the host properly. (See "Understanding the HBA Mask Field" on page 59 for more details.)

To register a host:

1. Choose **hosts > create host**.

   The **create host** window opens (Figure 23).
2. Enter the host's computer name in the **host name** field.
3. Enter the host's Fibre Channel worldwide name in the **host wwn** field.
4. Select the appliance target port to which the host HBA is connected in the **host target path** field. Each appliance node has six target ports, labeled 0 through 5 on the **register host** window.
5. Enter the host's operating system in the **host os** field.
6. Click **submit**.

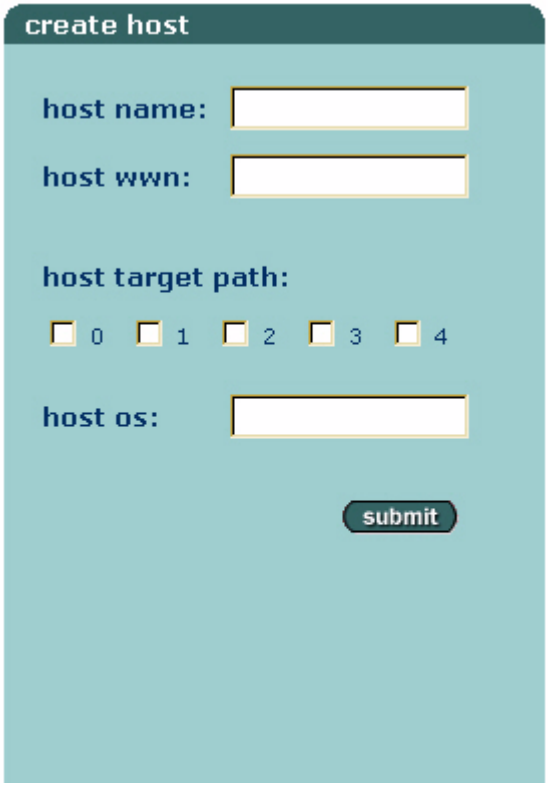

Figure 23: Create host window

Understanding the HBA Mask Field

Each appliance node has six target ports, numbered 0 through 5. The appliance software uses a bitmap to identify how hosts are connected to the target ports on each node. The bitmap results are translated into hexadecimal format to populate the **hba mask** field on the **view hosts** window.

For example, the host HBA and three appliance target ports (0, 1, and 2) are connected to a fabric switch zone. Figure 24 shows what the bitmap looks like when these three target ports are zoned with the host HBA.

Figure 24: Bitmap example

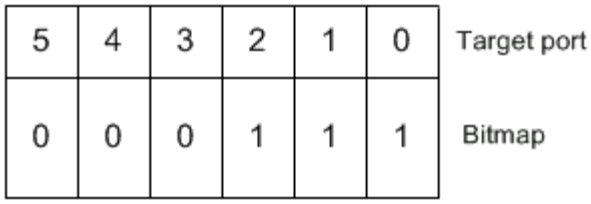

The hexadecimal translation of this bitmap is 0X7. This value displays in the **hba mask** field on the **view hosts** window.

Table 3 identifies some common HBA masks.

Table 3: Common HBA masks

Bitmap	Hexadecimal Translation	Host HBA to Target Port
0 0 0 0 0 1	0X1	Host HBA directly connected to target port 0
0 0 0 0 1 0	0X2	Host HBA directly connected to target port 1
0 0 0 1 0 0	0X4	Host HBA directly connected to target port 2
0 0 1 0 0 0	0X8	Host HBA directly connected to target port 3
0 1 0 0 0 0	0X10	Host HBA directly connected to target port 4
1 0 0 0 0 0	0X20	Host HBA directly connected to target port 5
0 0 0 1 1 1	0X7	Host HBA and target ports 0, 1, and 2 connected to fabric switch zone
1 1 0 0 0 0	0X30	Host HBA and target ports 4 and 5 connected to fabric switch zone
0 0 1 1 1 1	0X1F	Host HBA and target ports 0, 1, 2, 3, and 4 connected to fabric switch zone
1 1 1 1 1 1	0X3F	All target ports connected to fabric switch zone

Deleting Hosts

You may want to delete a host from the appliance configuration if:

- The host or one of its components is faulty.
- The host HBA has changed.
- The appliance will no longer communicate with the host.

Note: Only HP service personnel should delete hosts from the appliance configuration.

To delete a host:

1. Delete any LUN maps assigned to the host. (See "Deleting LUN Maps" on page 67.)
2. Choose **hosts > delete host**.

   The **delete hosts** window opens (Figure 25).
3. Choose the host you want to delete.
4. Click **submit**.

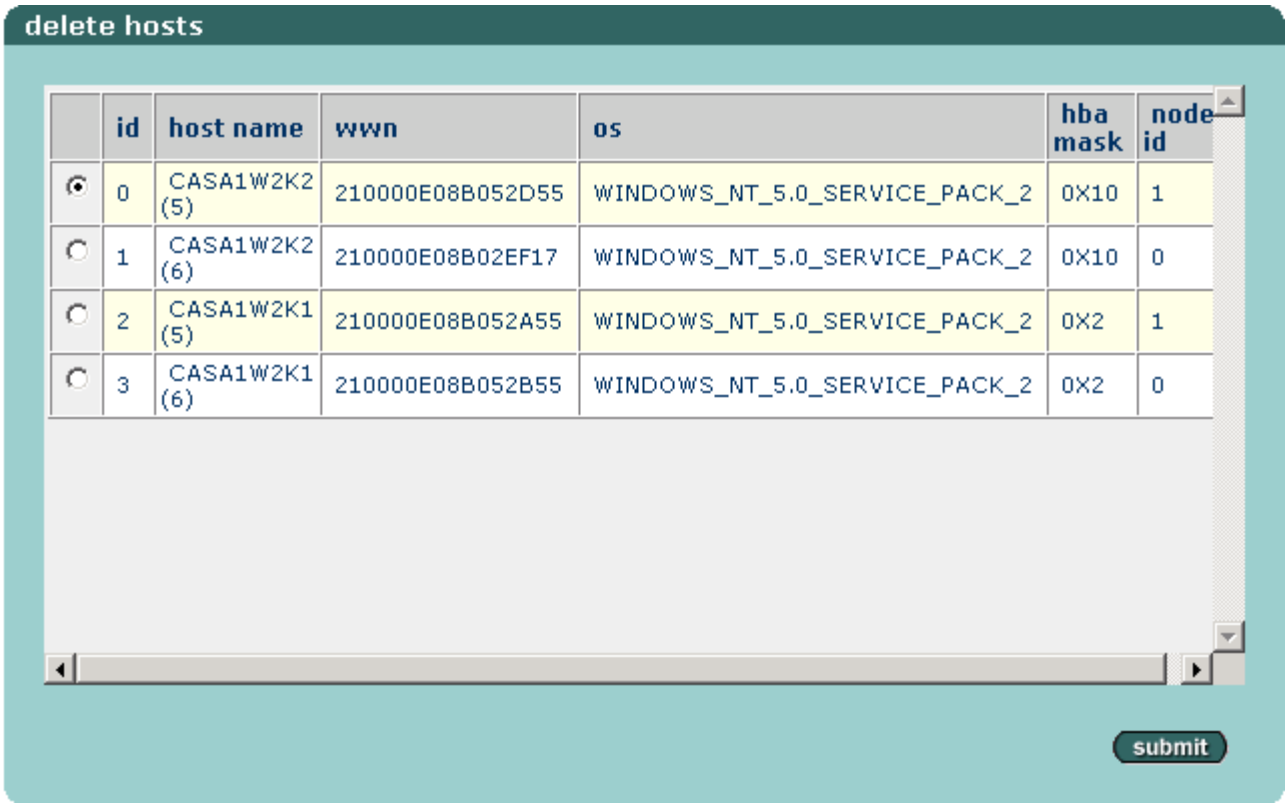

Figure 25: Delete hosts window

Managing LUNs

7

This chapter describes how you can use the appliance to manage logical unit numbers (LUNs). It covers the following topics:

- About LUNs, page 62
- Mapping LUNs, page 63
- Resizing LUNs, page 68
- Viewing LUN Information, page 78
- Viewing LUN History, page 81
- Viewing the LUN Tree, page 82
- Adding LUN Location, page 87
- Viewing LUN Statistics, page 90
- Viewing the LUN Matrix, page 92
- Deleting LUNs, page 94
- Rescanning LUNs, page 95

About LUNs

A logical unit number (LUN) identifies storage to a host. You can assign a LUN to one disk on a storage array or combine several disks as one LUN, creating the appearance of a single disk.

By virtualizing your storage, you can manage both physical LUNs (LUNs that exist) and virtual LUNs (LUNs that you create based on existing physical LUNs). The appliance supports a maximum of 256 physical LUNs and 4,096 virtual LUNs.

Mapping LUNs

When you map a LUN, you assign the LUN to a host HBA, which allows the host to access data on that LUN. You can map one LUN to one host or to several hosts. Mapping LUNs through the appliance increases security because you assign access rights from the network, not directly from the host. Another host connected to the appliance cannot access a LUN or another host's data unless you map the LUN to that host.

Because hosts can share LUNs, you must choose carefully which hosts have access to the same LUNs. For example, you assign two hosts with different operating systems to the same LUN. Either host may overwrite data on that LUN.

This section covers the following topics:

- About LUN Maps
- LUN Map Rules
- Creating LUN Maps
- Deleting LUN Maps

About LUN Maps

When you create a LUN map, you must specify the following information:

- **Host**—The host HBA you want to map to a LUN.
- **LUN**—The number identifying the storage disk you want to map to a host.
- **Host LUN**—The number identifying the LUN to the host. You can specify a value from 0 to 255 for each target port.
- **Target**—The target port from the appliance node that is connected to the host. You identify which target ports are connected to a host when you register the host. (See "Registering Host Information" on page 55 for more information.)
- **Access control**—The type of access the host has to the LUN: **read-only** or **read write**.

LUN Map Rules

Once you map a LUN, you cannot make changes to it. For example, you cannot use a mapped LUN:

- To create a LUN partition
- To create a LUN expansion
- As the target LUN of either a local FCP mirror or an IP mirror
- As a point-in-time image target

You can create a LUN partition or a LUN expansion and then map that LUN to a host.

Creating LUN Maps

To create a LUN map:

1. Choose **LUN Map > create lunmap**.

   The **create lunmap** window opens (Figure 26).
2. Choose the host HBA you want to use in the **select host** field.
3. Choose the LUN you want to map in the **select lun** field.
4. Enter the number that identifies the LUN to the host in the **select host lun number** field.

   You can re-use the same host LUN for another LUN map between the same host and a different LUN, but you cannot use the same appliance target port.
5. Choose **read only** or **read-write** in the **access control** field.
6. Choose an appliance target port in the **select target** field.
7. Click **submit**.

   The **view lunmaps** window opens, displaying the new LUN map (Figure 27).

Note: The **select host lun number** field on the **create lunmap** window and the **visible lun** field on the **view lunmaps** window reflect the same LUN, that is, the number that the host uses to identify the LUN.

Repeat steps 1 through 7 to create another path to the host for the LUN if *all* of the following criteria apply:

- There are two paths to the appliance, one on Node 0 and one on Node 1.
- You are using failover software on the host.
- The host has two or more HBAs.

Note: If you are not using failover software and the host has only one HBA, do not repeat this procedure on Node 1.

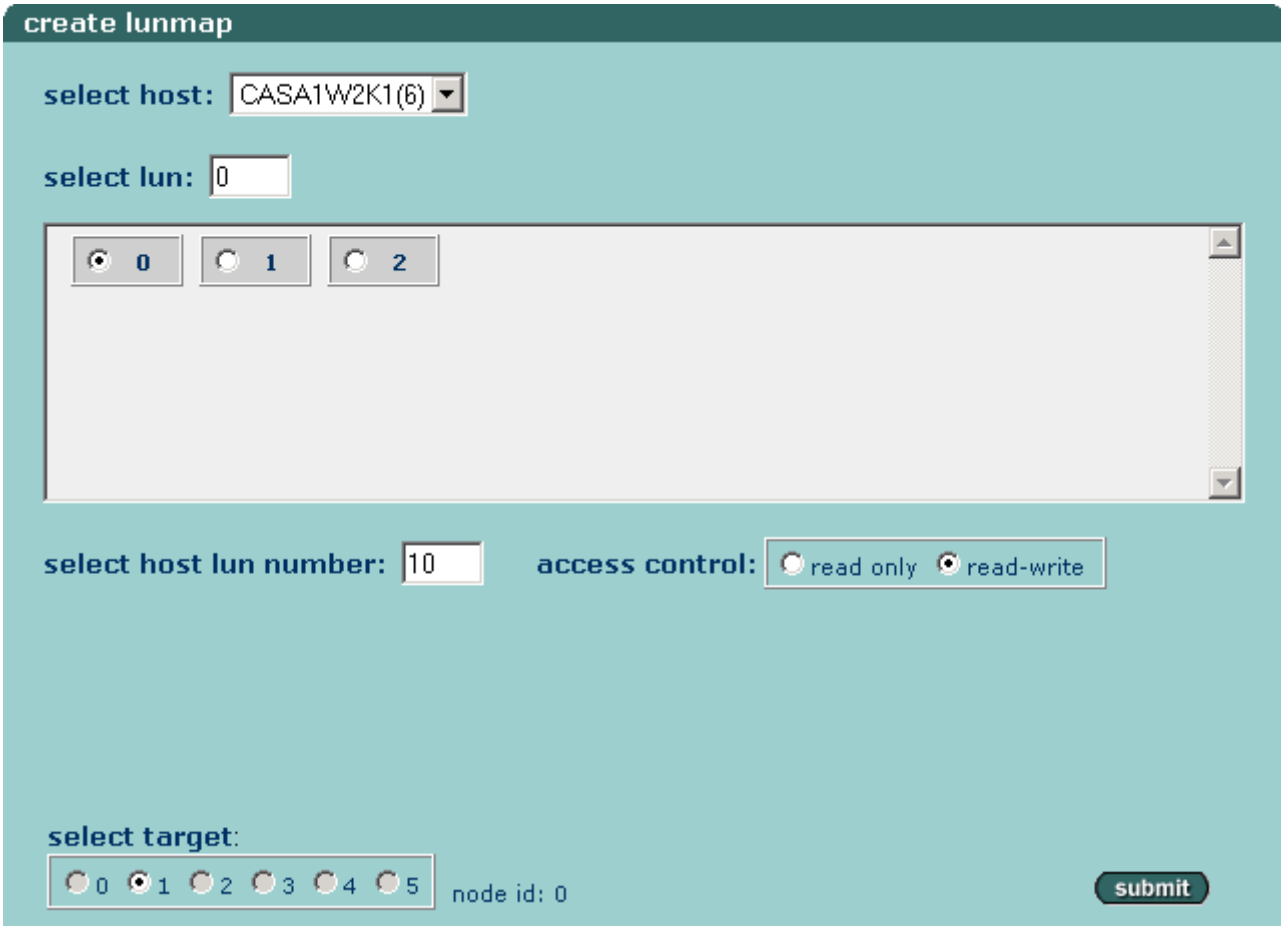

Figure 26: Create lunmap window

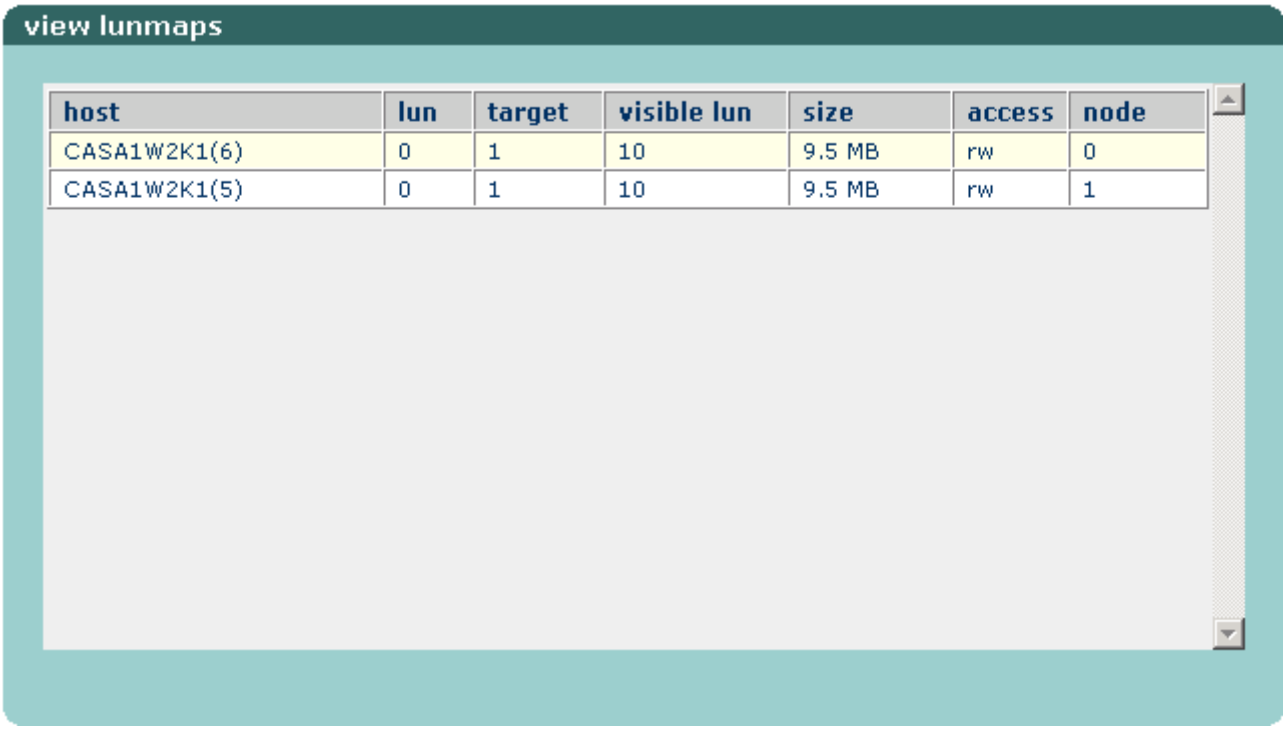

Figure 27: View lunmaps window

Deleting LUN Maps

You may need to delete a mapped LUN if:

- You selected the wrong host.
- You no longer want a host to access a LUN.
- You want to delete a LUN partition or LUN expansion that is mapped to a host. You must delete the LUN map before you can delete the partition or expansion.

After you delete a LUN, it is no longer visible to the host.

To delete a LUN map:

1. Choose **LUN Map > delete lunmaps**.

   The **delete lunmaps** window opens (Figure 28).
2. Choose the LUN map you want to delete.
3. Click **submit**.

   The **view lunmaps** window opens. The LUN map you deleted is not listed.

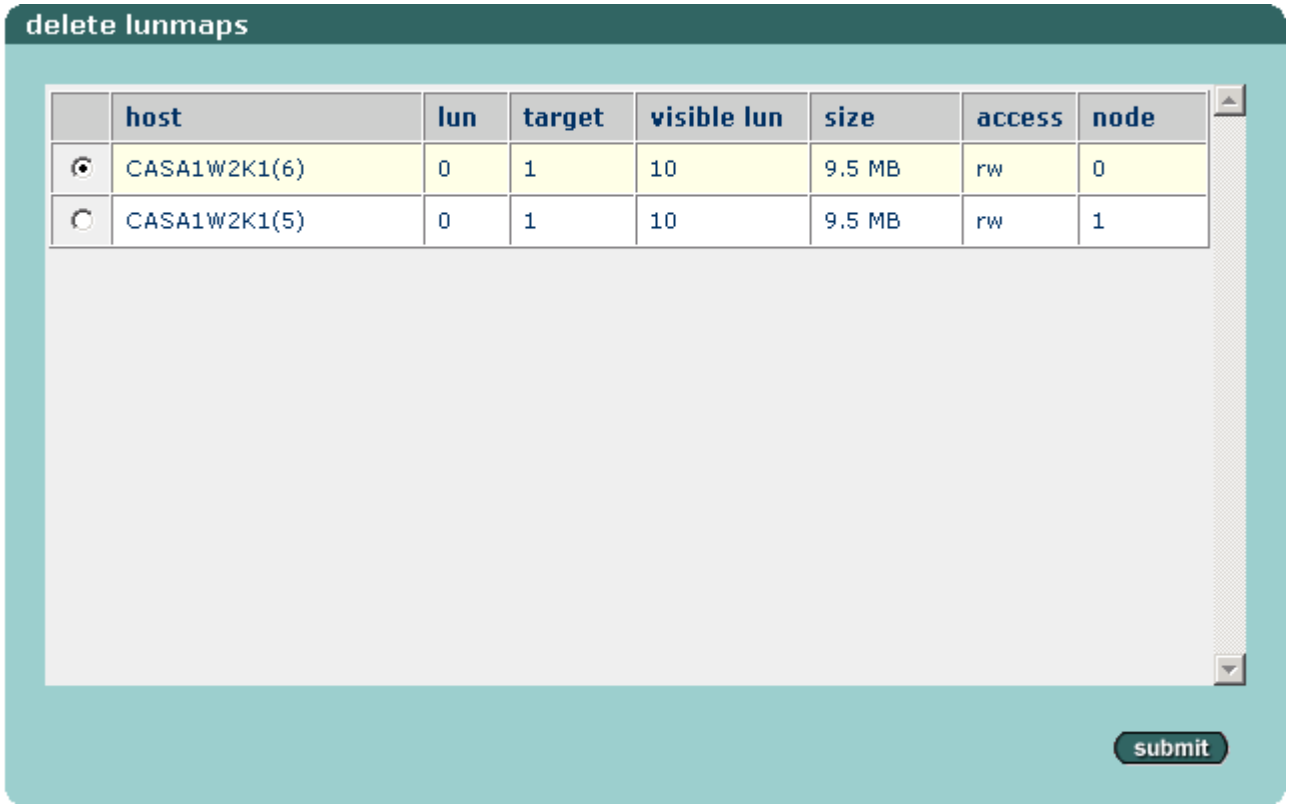

Figure 28: Delete lunmaps window

Resizing LUNs

You can partition LUNs (divide one LUN into one or more smaller LUNs) or expand LUNs (combine multiple LUNs into one LUN) to create pooled storage areas that make it easier to store and manage information.

Before you resize a LUN, HP recommends that you back up the data on the LUN.

This section covers the following topics:

- Viewing Resized LUNs
- Creating LUN Partitions
- Deleting LUN Partitions
- Creating LUN Expansions
- Deleting LUN Expansions

Viewing Resized LUNs

To view all partitioned and expanded LUNs, choose **LUN resize** from the menu bar. The **view resizes** window opens (Figure 29), displaying the following information:

- **type**—The type of LUN resize (partition or expansion).
- **original lun**—The number of the physical LUN from which the partition or expansion originated.
- **other lun**—For a partition, this number identifies the virtual LUN of the partition that you created. For an expansion, this number identifies the LUN that you added to the original LUN to create the expanded LUN.
- **size (blocks)**—The size of the partition or expansion in blocks.

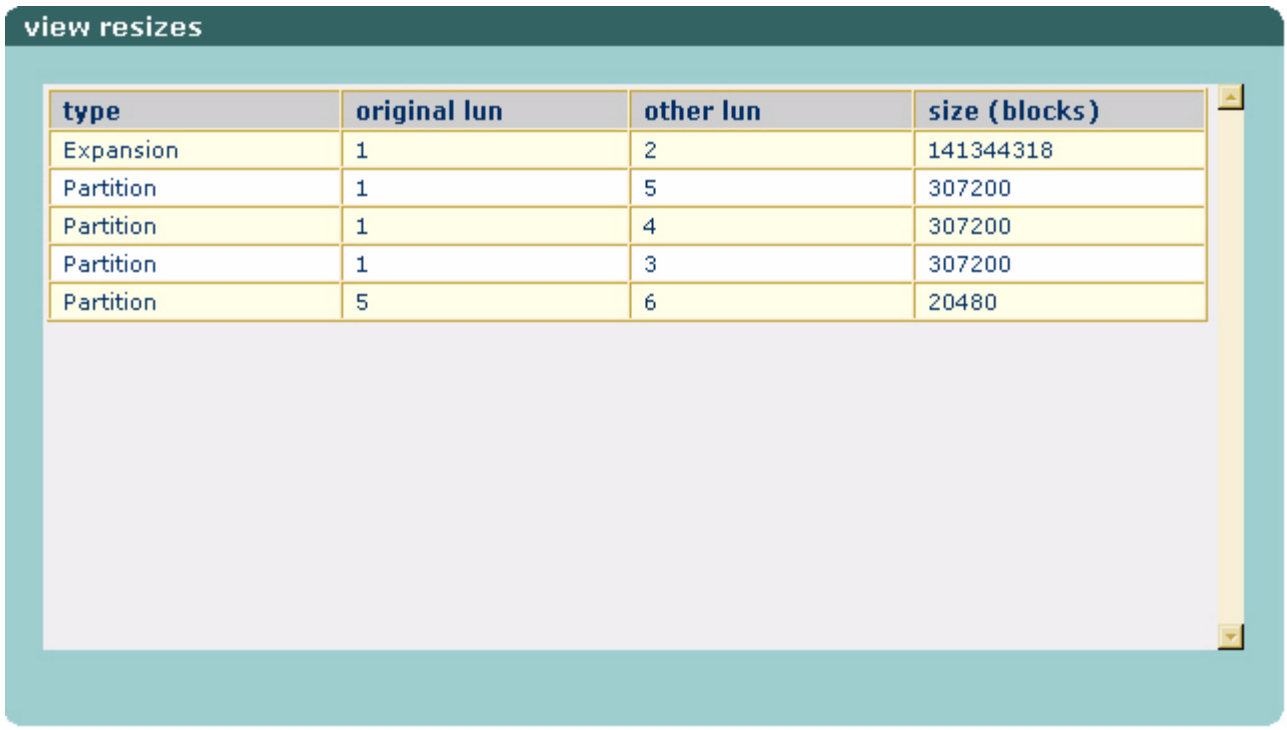

Figure 29: View resizes window

Creating LUN Partitions

LUN Partition Rules

Before you partition a LUN, ensure that the LUN is not:

- An exported LUN
- Part of a LUN expansion
- A mirror source or target
- A mapped LUN
- A point-in-time image source or target

The appliance software assigns a number to the partition you create. The number is based on the last available LUN and is incremented by one. In Figure 30, the last available LUN is 11. If you create a partition from LUN 6, the appliance software assigns LUN 12 to the partition (Figure 31).

Procedure

To create a LUN partition:

1. Choose **LUN resize > create partition**.

   The **create partition** window opens (Figure 30).
2. Choose the LUN you want to partition in the **select lun** field.
3. Choose the size of the partition in the **partition size** field.
4. Choose the number of partitions you want to create in the **number of partitions** field.

   You can create a maximum of three partitions when using the user interface. If you use the CLI, you can create an unlimited number of partitions. Refer to "Partition and Expansion Commands" in the *HP OpenView Continuous Access Storage Appliance Command Line Reference Guide* for more information.
5. Click **submit**.

   The **view partitions** window opens, displaying the partition you created (Figure 31).

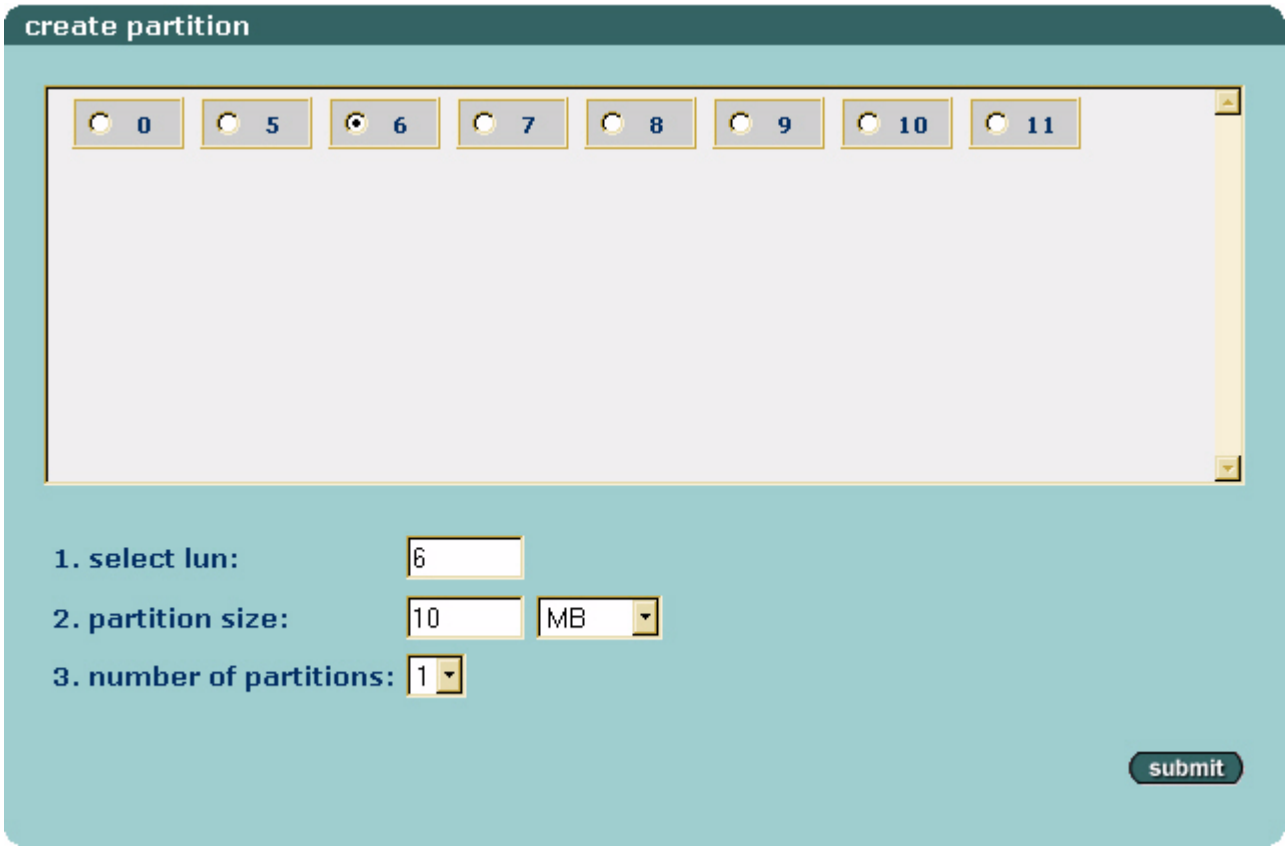

Figure 30: Create partition window

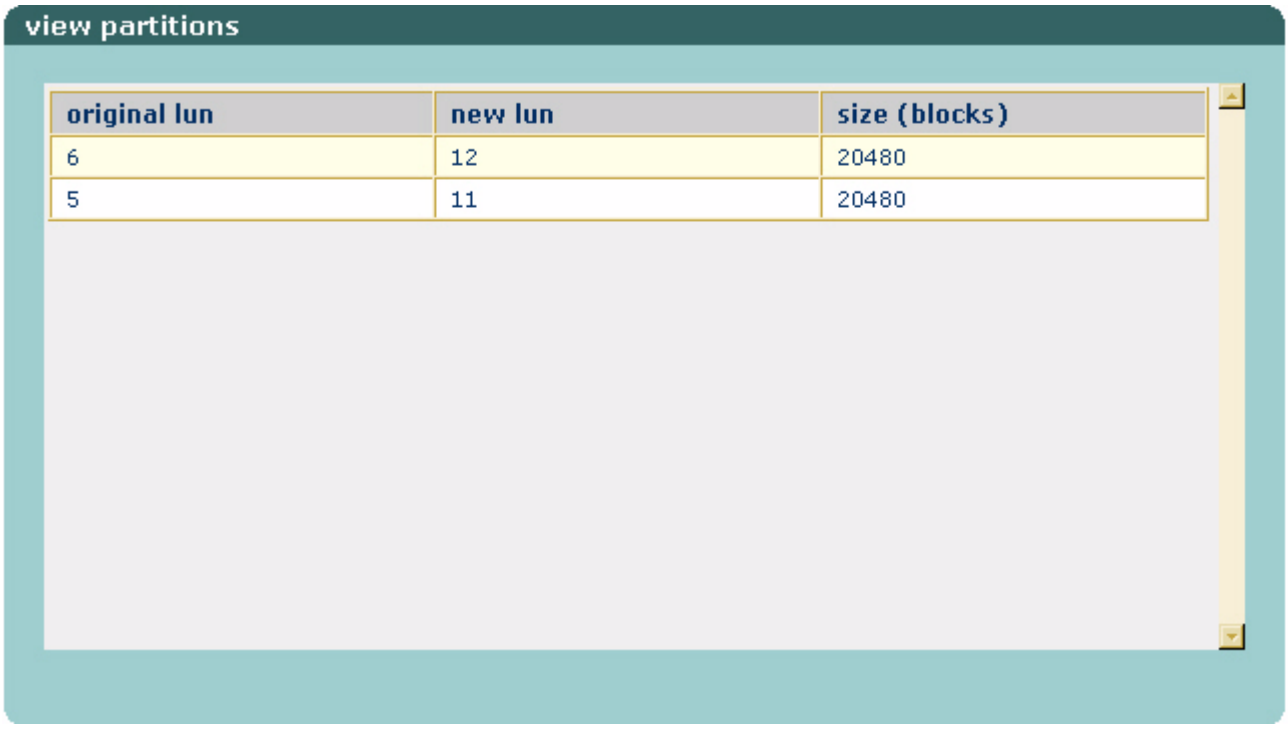

Figure 31: View partitions window

Deleting LUN Partitions

When you delete a LUN partition, you do it in “last created, first deleted” order.

To delete a LUN partition:

1. Choose **LUN resize > delete partition**.

   The **delete partitions** window opens (Figure 32).

2. Choose the partition you want to delete and click **submit**.

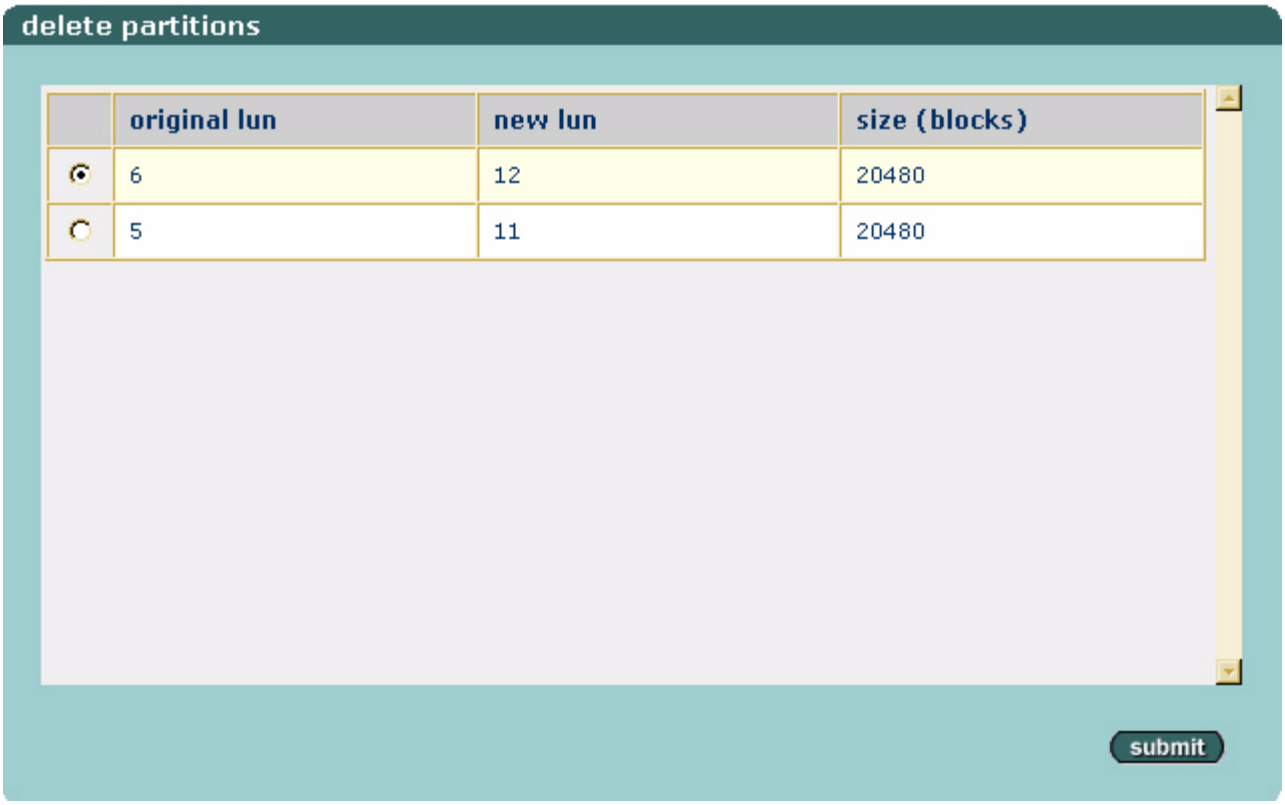

Figure 32: Delete partitions window

Creating LUN Expansions

About LUN Expansion

When you expand a LUN, you combine multiple LUNs to create one virtual LUN. You can expand a LUN when it reaches full capacity and requires more space.

A LUN expansion consists of an *expanded LUN* and a *merged LUN*. The merged LUN is added to the expanded LUN. In Figure 33, LUN 1 is the merged LUN and LUN 0 is the expanded LUN.

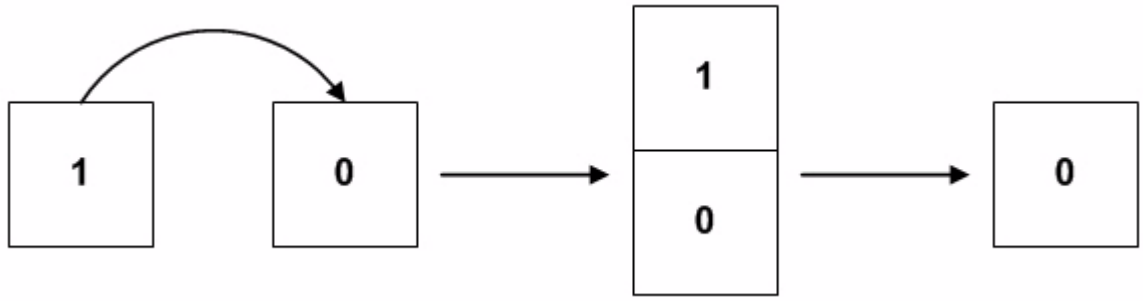

Figure 33: Expanding a LUN

Once a merged LUN is used in an expansion, it cannot be used for anything else. In Figure 33, LUN 1 is the merged LUN. Once it is combined with LUN 0, LUN 1 is unavailable; it ceases to exist as LUN 1.

However, the expanded LUN (LUN 0) can become a merged LUN to create another LUN expansion (Figure 34).

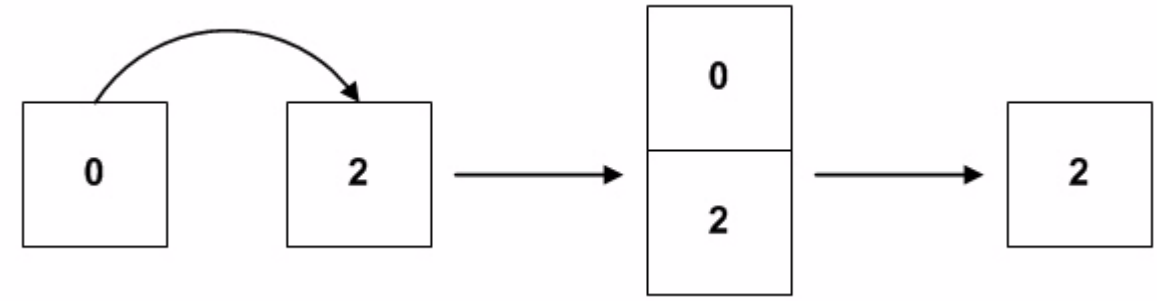

Figure 34: Creating a LUN expansion using an expanded LUN

LUN Expansion Rules

Before you expand a LUN, ensure that the LUN is not:

- An exported LUN
- A mirror source or target
- The merged LUN has not been mapped
- A point-in-time image source or target

LUN Expansion Types

There are two types of LUN expansions:

- **Virtual**—Expanding a LUN using the appliance software, which is described in this chapter. You are not limited by the types of storage arrays connected to the appliance.
- **Hardware**—Expanding a LUN from the storage array (such as the EVA). To successfully create a hardware LUN expansion, you must unmap (if the LUN is mapped) and delete the LUN from the appliance. Complete the expansion of the LUN from the storage array and perform a LUN rescan so the appliance can identify the new LUN. (See "Rescanning LUNs" on page 95 for more information.)

Note: Currently, the appliance only supports virtual LUN expansions.

Procedure

To create a LUN expansion:

1. Choose **LUN resize > create expansion**.

   The **create expansion** window opens (Figure 35).
2. Choose the expanded LUN in the **select LUN to expand** field.
3. Choose the merged LUN in the **select LUN to expand with** field.
4. Click **submit**.

   The **view expansions** window opens, displaying the new LUN expansion (Figure 36). The newest expansion is listed first; the oldest expansion is listed last.

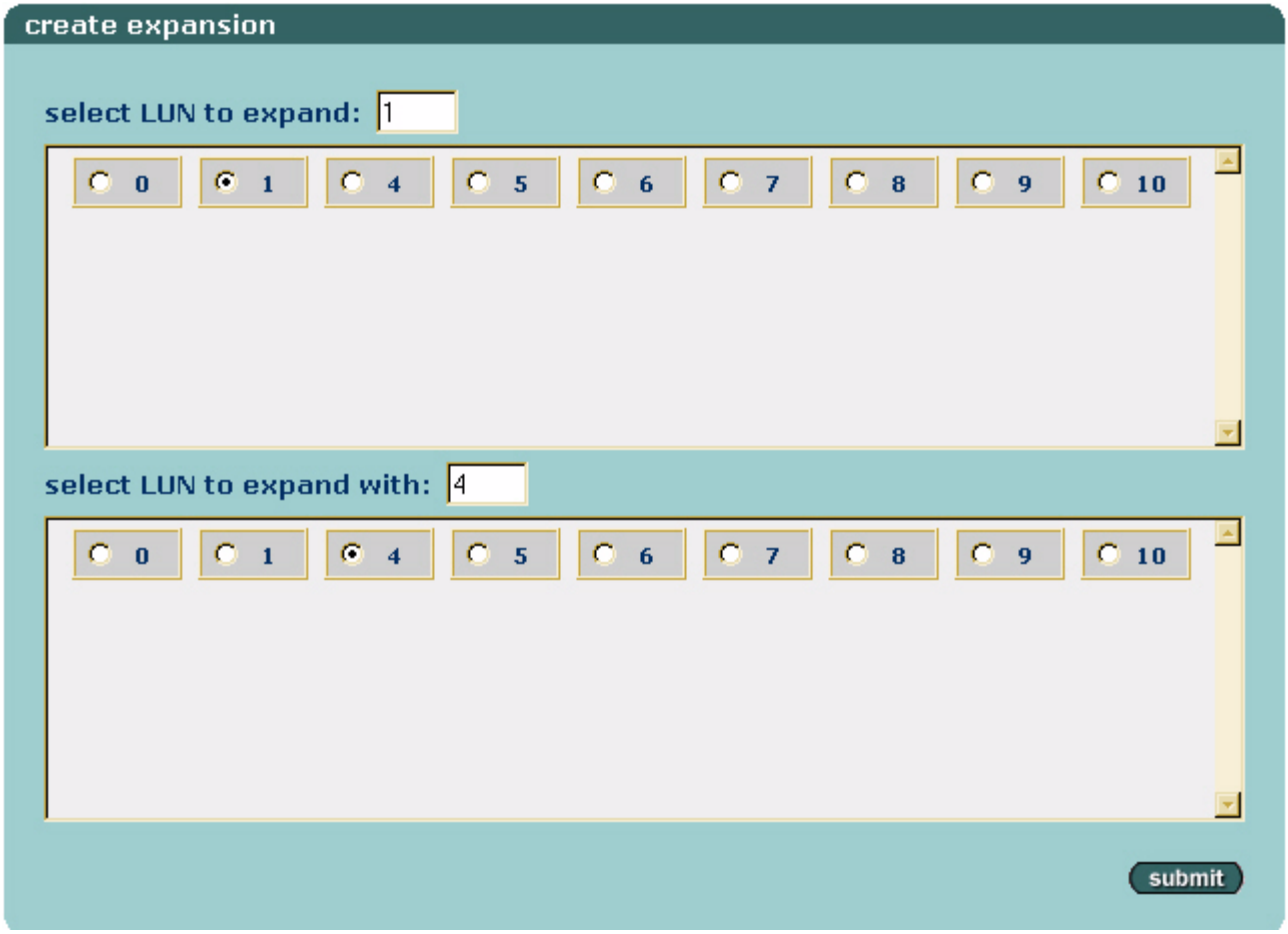

Figure 35: Create expansion window

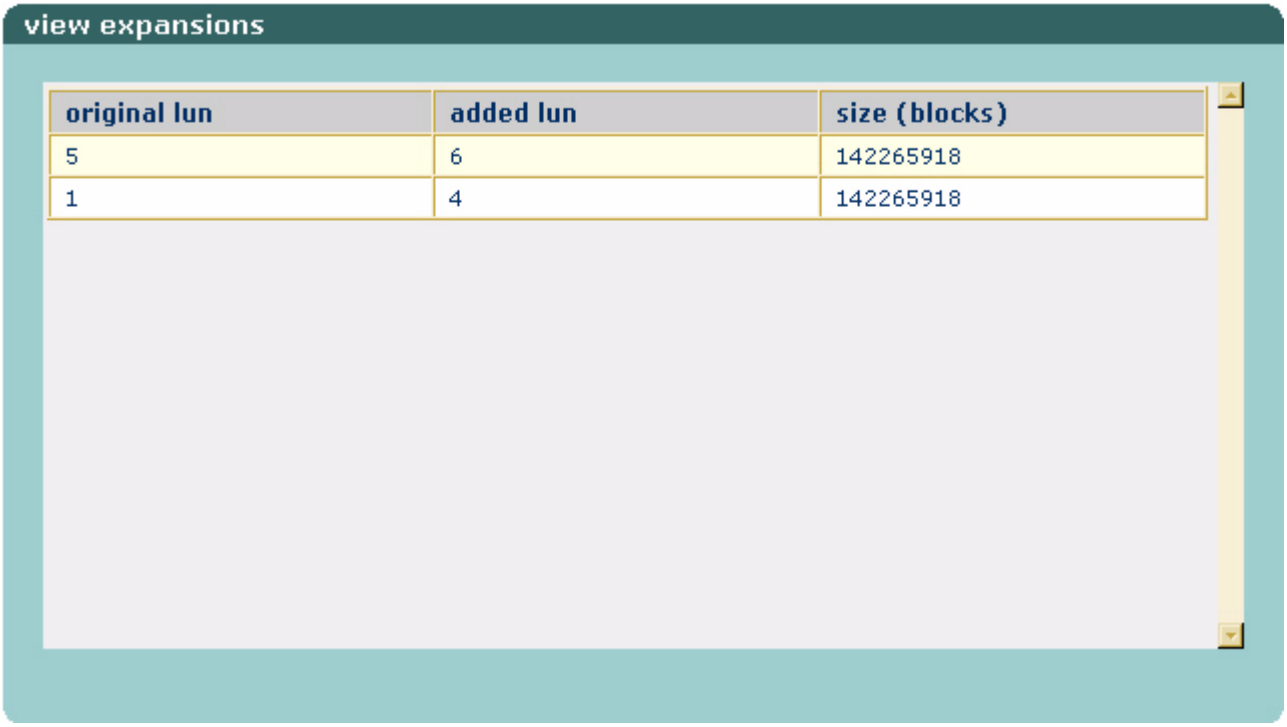

Figure 36: View expansions window

Deleting LUN Expansions

When you delete a LUN expansion, you do it in “last created, first deleted” order.

To delete a LUN expansion:

1. Choose **LUN resize > delete expansion**.

   The **delete expansions** window opens (Figure 37).
2. Choose the expansion you want to delete.
3. Click **submit**.

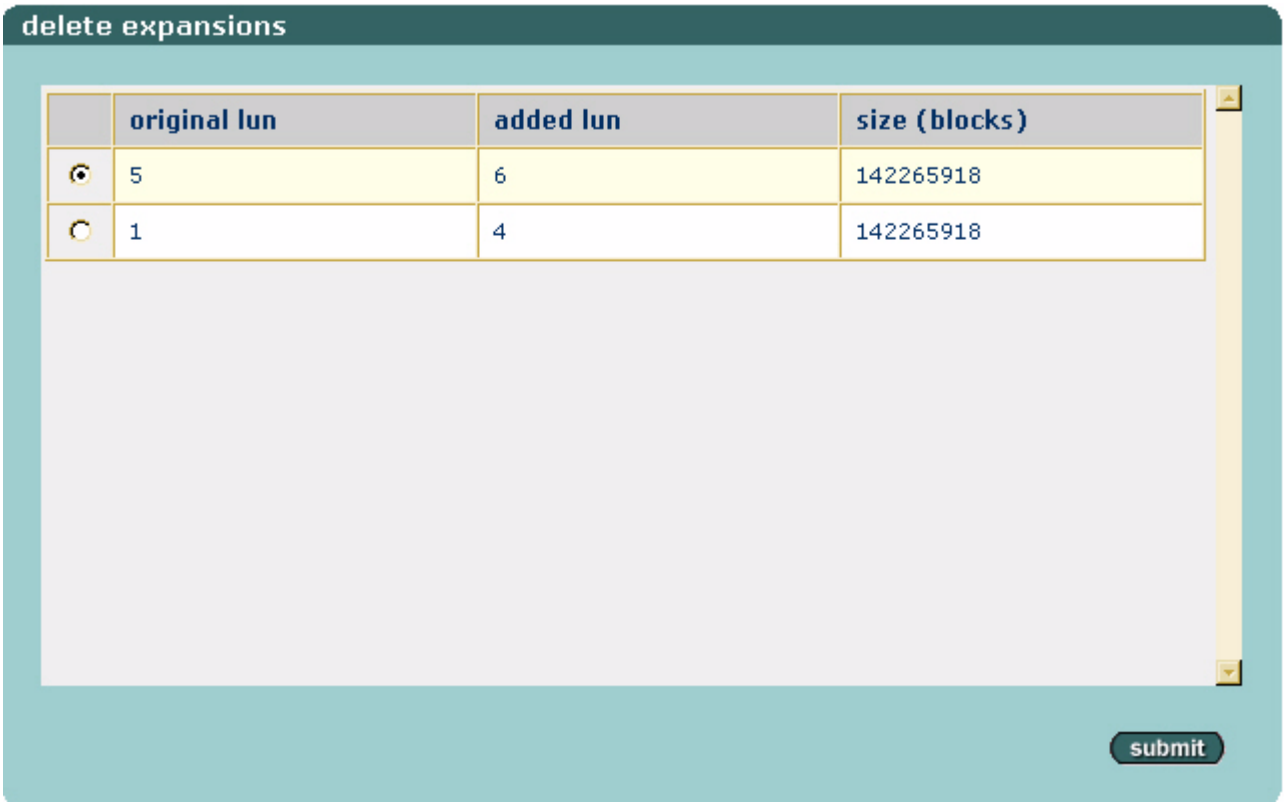

Figure 37: Delete expansions window

Viewing LUN Information

You use the **view luns** window to view LUN information such as the size and storage vendor.

To view LUN information:

1. Choose **LUNs** from the menu bar.

   The **view luns** window opens (Figure 38).

2. Select **details** for a LUN.

   The **lun x window** opens (Figure 39), displaying detailed information about the LUN. Scroll down to view all details.

Navigating the view luns window

The **viewing luns x - x** field in the lower right corner of the **view luns** window indicates which LUNs are currently displayed. Because the window only shows four LUNs within the space allotted, use the scroll bar to scroll and view other LUNs.

To view a specific range of LUNs:

1. Enter the number of the LUN you want to begin with in the **start lun** field.
2. Choose the quantity of LUNs you want to view in the **show next** field. You can choose 5, 10, 15, 20, or 25; the default is 10.
3. Click **GO**.

   The **view luns** window displays the LUNs based on your selections.
   The **viewing luns x - x** field value changes to reflect the range of LUN s you are viewing. For example, if you enter 11 in the start lun field and choose 20 in the show next field, the view luns window displays LUNs 11 through 30.

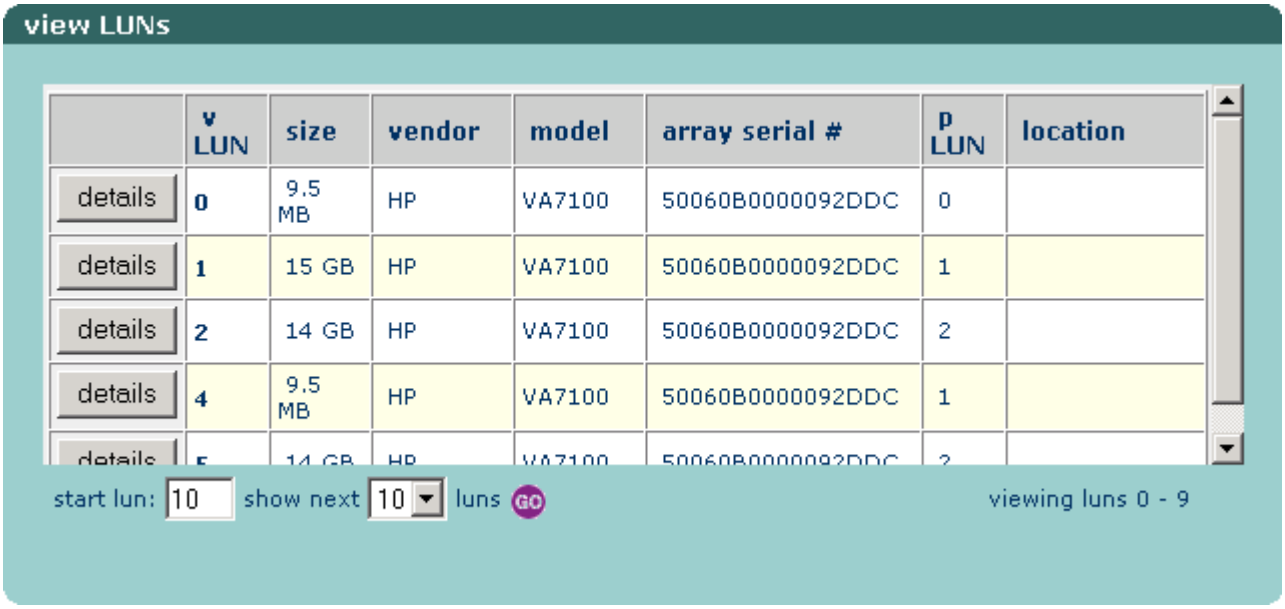

Figure 38: View luns window

Figure 39: lun x window

Table 4 describes the fields on the **lun x** window.

Table 4: Fields on the lun x window

Field	Description
size	Size, in gigabytes, of the LUN.
inquiry	Data returned by the SCSI inquiry command, indicating the device vendor name and serial ID.
lun serial	ID assigned when the LUN was created.
blocks	Number of 512-byte blocks on the LUN.
vendor name	Storage vendor of the physical storage device.
health	Health of the LUN. UP indicates that the LUN is connected and working properly. DOWN indicates that the LUN is not physically connected to the appliance and is not functioning.
hba	The host bus adapter on the appliance to which the LUN is connected. The appliance uses the values of the **hba**, **scsi id**, and **bus** fields to identify the LUN connected to the appliance.
sector size	Physical space on the LUN. This field value is always 512.
physical lun	Number of the physical (or native) LUN. If you are viewing a virtual LUN, the physical LUN is a different number.
scsi id	The appliance uses the values of the **hba**, **scsi id**, and **bus** fields to identify the LUN connected to the appliance.
bus	The appliance uses the values of the **hba**, **scsi id**, and **bus** fields to identify the LUN connected to the appliance.

The **lun x** window also indicates how you are using the LUN by displaying **true** or **false** in the following fields:

- **partition**
- **fcp mirror source**
- **fcp mirror target**
- **remote mirror source**
- **remote mirror target**
- **partitioned**
- **expanded**
- **mapped**
- **pt image source**
- **ip mirror target**
- **hp service guard**
- **imported**
- **exported**

Viewing LUN History

To view the history of a LUN:

1. Select **history** from the **lun x** window (Figure 39).
   The **history of lun x** window opens, displaying the history of the LUN (Figure 40).
2. Choose a LUN from the **lun** column to view the history of that LUN.

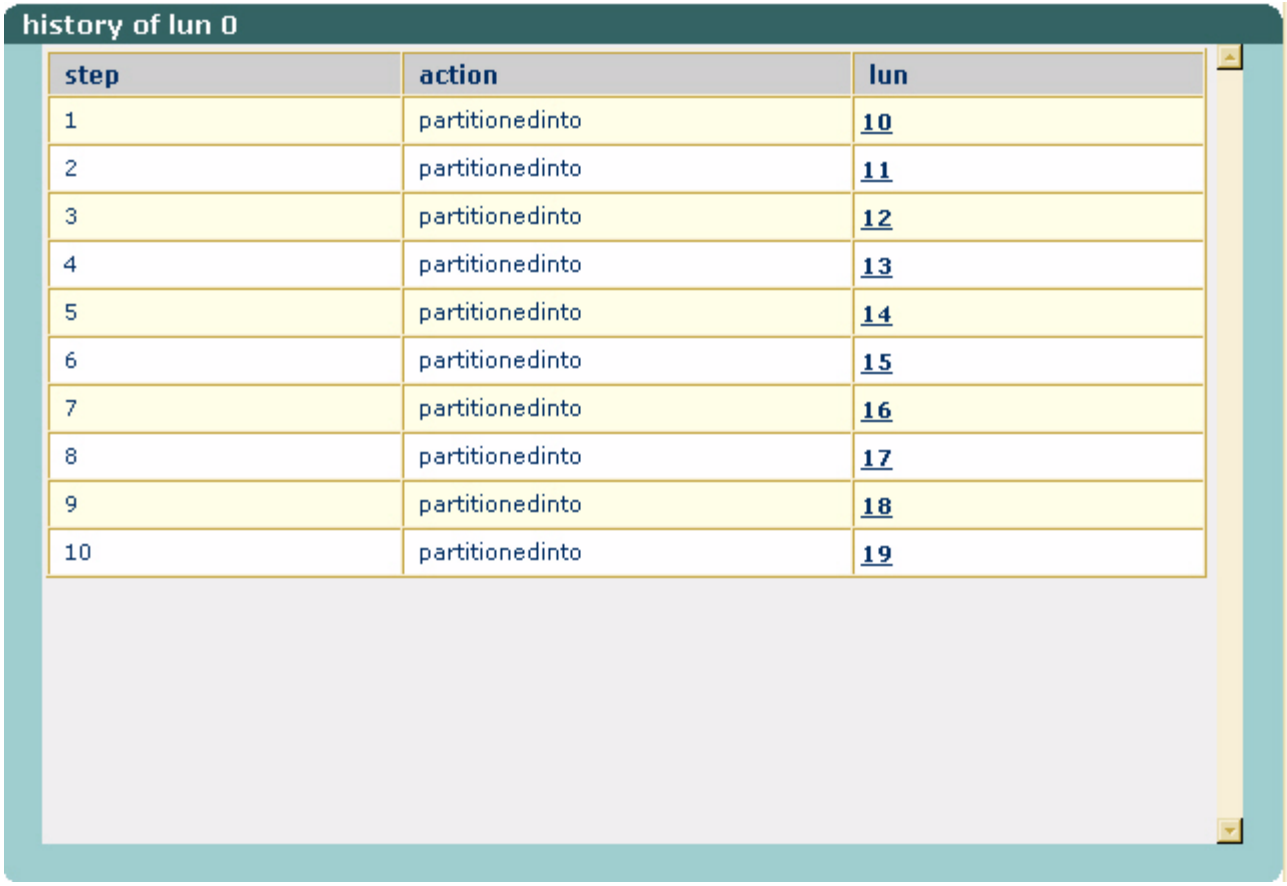

Figure 40: History of lun x window

Viewing the LUN Tree

The LUN tree enables you to view physical LUNs and virtual LUNs separately. The LUN tree presents LUN information in a hierarchical format. For physical LUNs, the hierarchy begins with the storage vendor. From the storage vendor, you can drill down through the storage model and storage serial number to a physical LUN. For virtual LUNs, the hierarchy begins with a range of virtual LUNs. The size of the range depends on how many virtual LUNs the appliance software is managing. From the range, you can drill down to individual virtual LUNs.

Viewing Physical LUNs

To view a physical LUN:

1. Choose **LUNs > LUN tree > storage**.

   In the **LUN tree** submenu, the **storage** hierarchy expands so you can view the storage vendors. The **LUN tree: storage** window displays a prompt to select a specific vendor (Figure 41).

2. Select a storage vendor in the **LUN tree** submenu.

   A list of storage models displays below the storage vendor you selected.
   The **LUN tree : storage** window displays a prompt to select a specific model.

3. Select a storage model in the **LUN tree** submenu.

   A list of serial numbers displays below the storage model you selected.
   The **LUN tree : storage** window displays a prompt to select a specific serial number.

4. Select a serial number in the **LUN tree** submenu.

   A list of physical LUNs displays below the serial number you selected.
   The **LUN tree : storage** window displays a prompt to select a specific physical LUN (Figure 42).

5. Select a physical LUN in the **LUN tree** submenu.

   The v**iew LUN tree** window displays information about the LUN you selected (Figure 43). Figure 43 shows information about LUN 1.

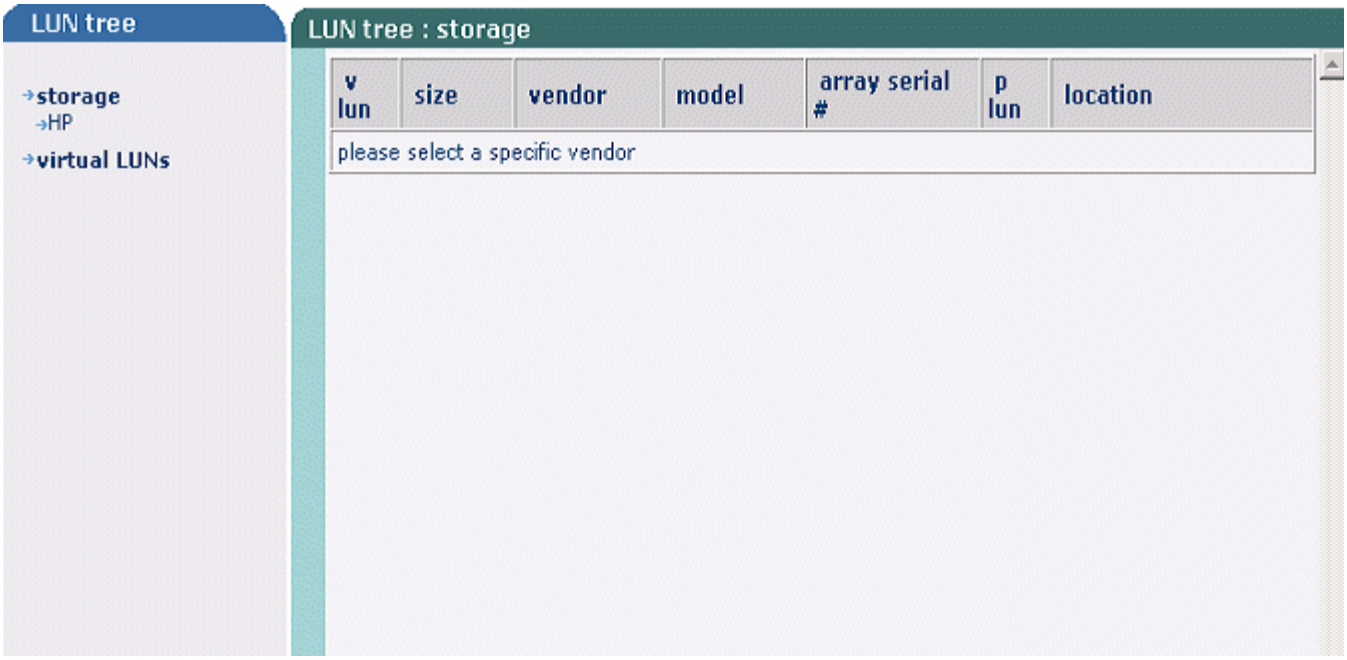

Figure 41: LUN tree : storage window (selecting a storage vendor)

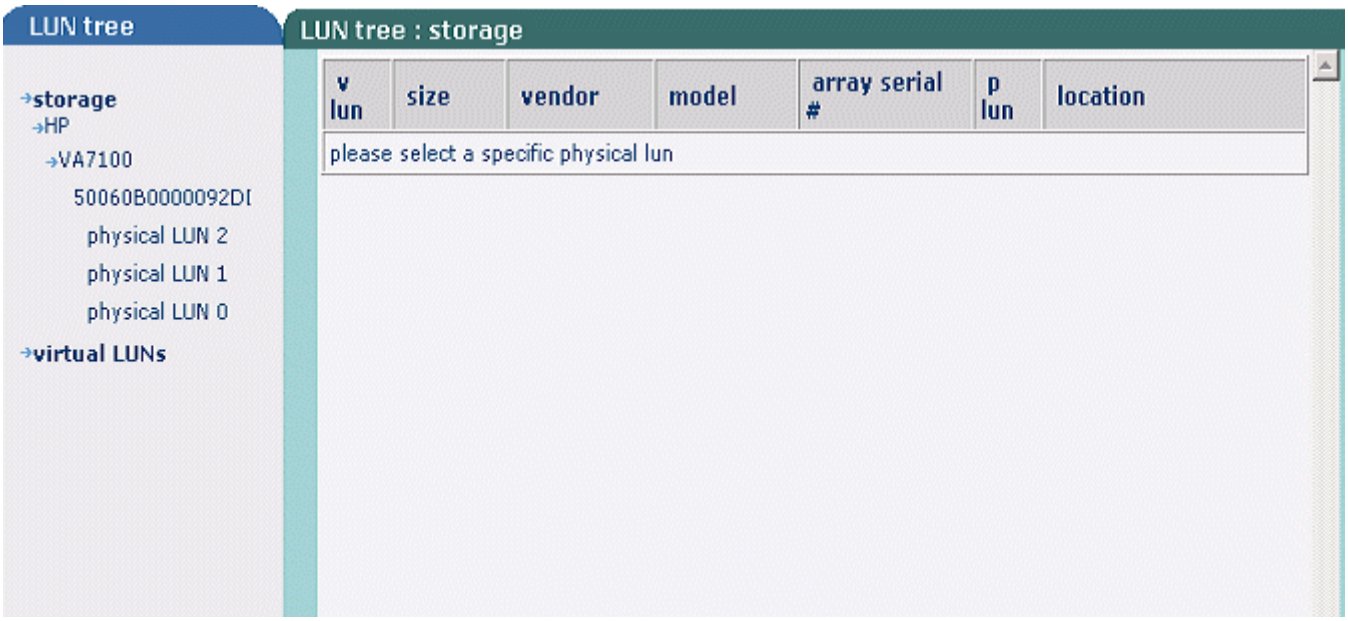

Figure 42: LUN tree : storage window (selecting a physical LUN)

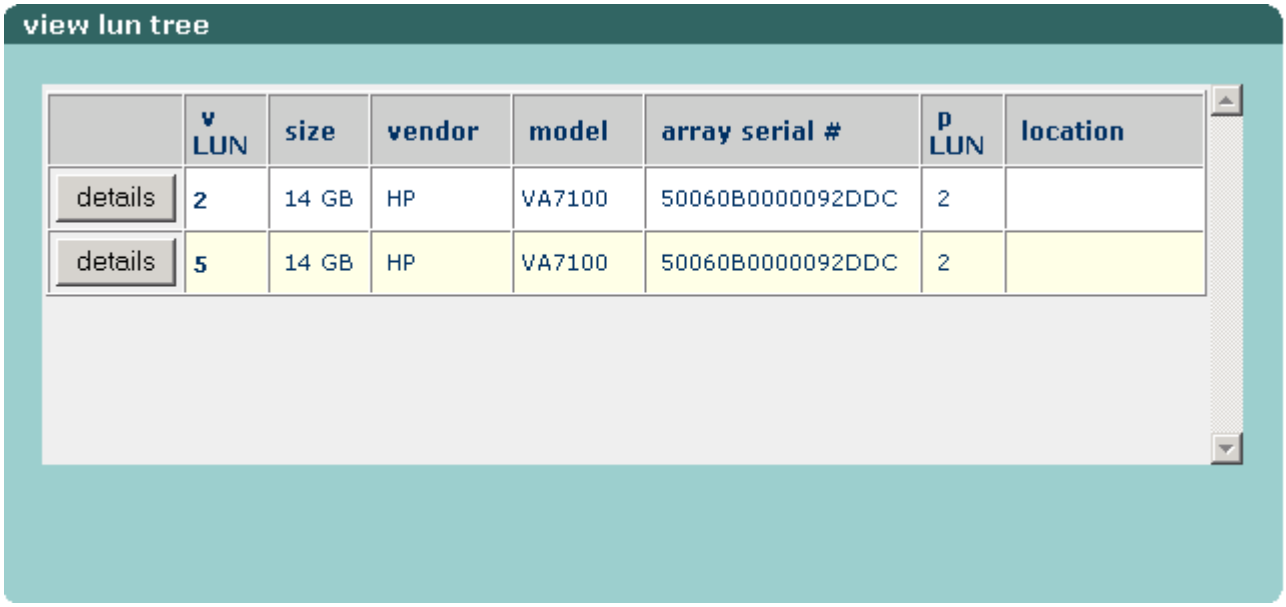

Figure 43: View lun tree window

Viewing Virtual LUNs

To view a virtual LUN:

1. Choose **LUNs > lun tree > virtual luns**.

   In the **LUN tree** submenu, the **virtual luns** hierarchy expands so you can view the virtual LUN ranges. The **LUN tree: virtual luns** window displays a prompt to select a specific virtual LUN range (Figure 44).

2. Select a virtual LUN range in the **LUN tree** submenu.

   A list of virtual LUNs displays below the range you selected. The **LUN tree : virtual luns** window displays a prompt to select a specific virtual LUN (Figure 45).

3. Select a virtual LUN in the **LUN tree** submenu.

   The **view lun tree** window displays information about the LUN you selected (Figure 46).

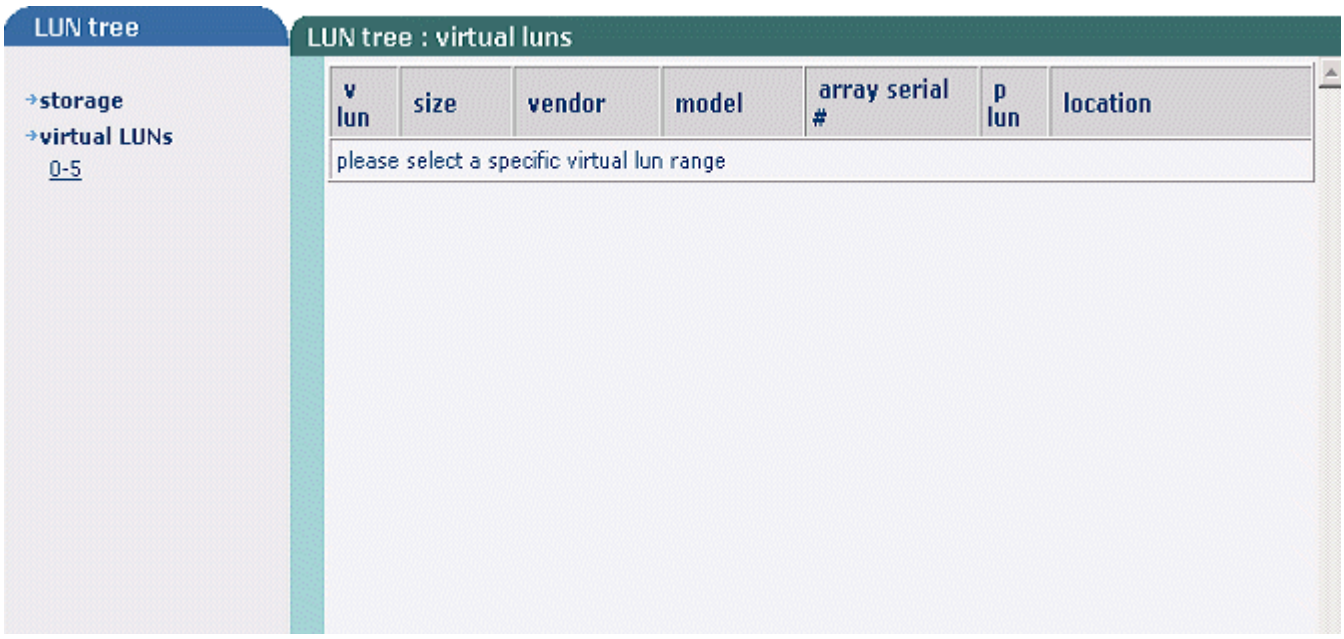

Figure 44: LUN tree : virtual luns window (selecting a LUN range)

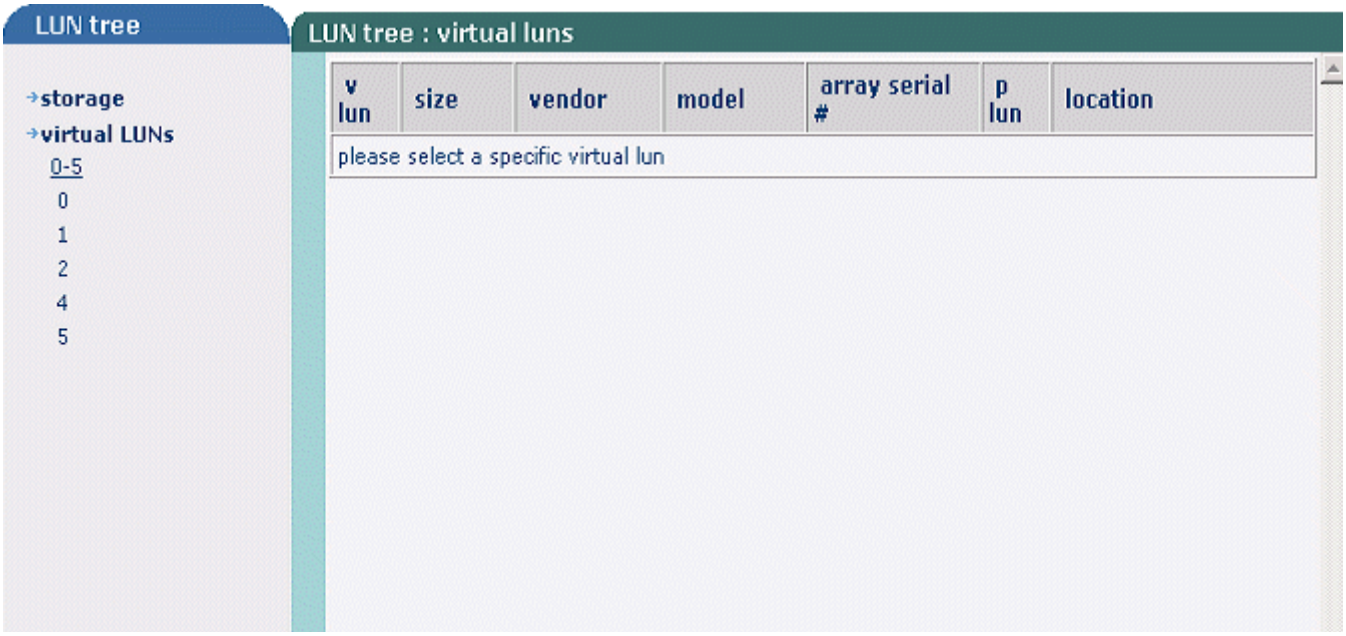

Figure 45: LUN tree : virtual luns window (selecting a virtual LUN)

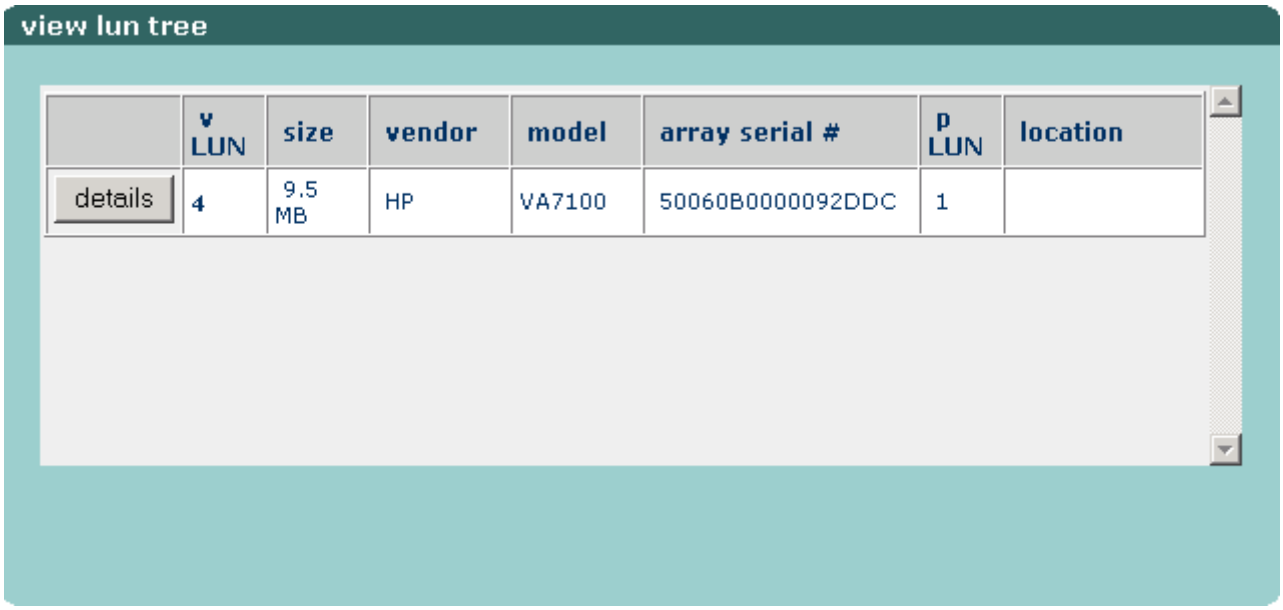

Figure 46: View lun tree window

Adding LUN Location

You can add the location for the LUNs that you manage through the appliance.

Note: You can only add a location for a physical LUN.

To add the location for a physical LUN:

1. Choose **LUNs > set lun location**.

   The **set LUN location** window opens (Figure 47).
2. Do either of the following to choose a LUN:
   - Select the checkbox for a LUN.
   - Enter the LUN in the **select LUN(s)** field.
3. Enter the location in the **input location** field. To include spaces, enclose the text in quotation marks.
4. Click **submit**.

   The location displays in the **location** field of the **view luns** window (Figure 48) and in the **location** field of the **view lun tree** window (Figure).

To delete the location for a physical LUN:

1. Choose **LUNs > set lun location**.

   The **set LUN location** window opens (Figure 47).
2. Do either of the following to choose a LUN:
   - Select the checkbox for a LUN.
   - Enter the LUN in the **select LUN(s)** field.
3. Delete the text from the **input location** field.
4. Click **submit**.

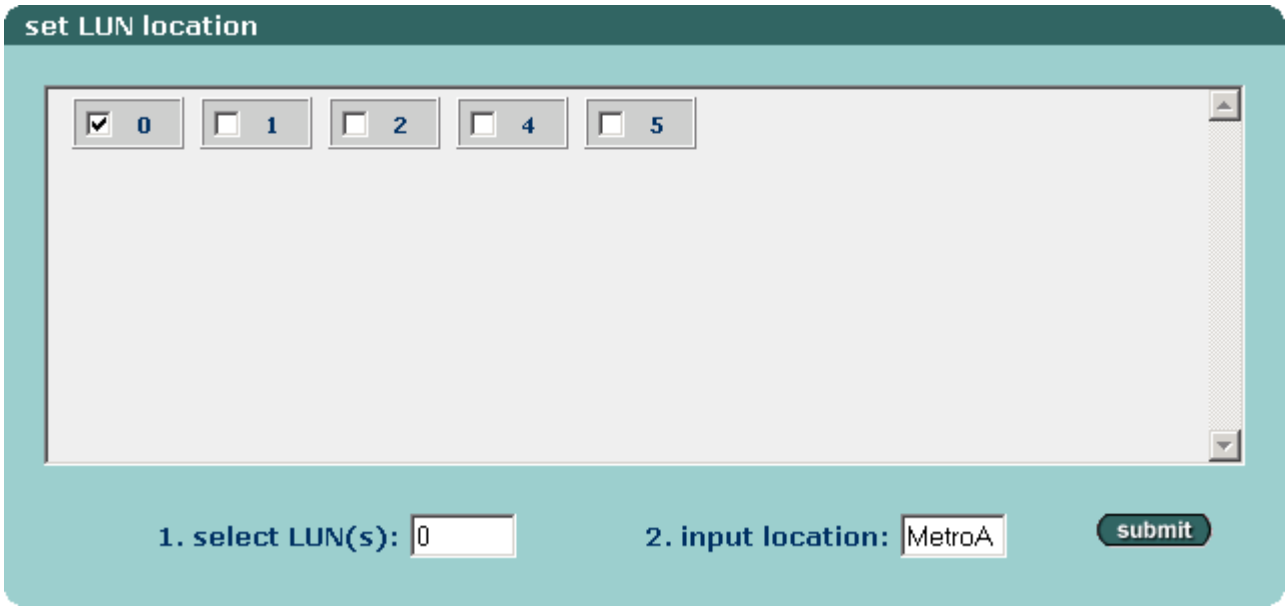

Figure 47: Set LUN location window

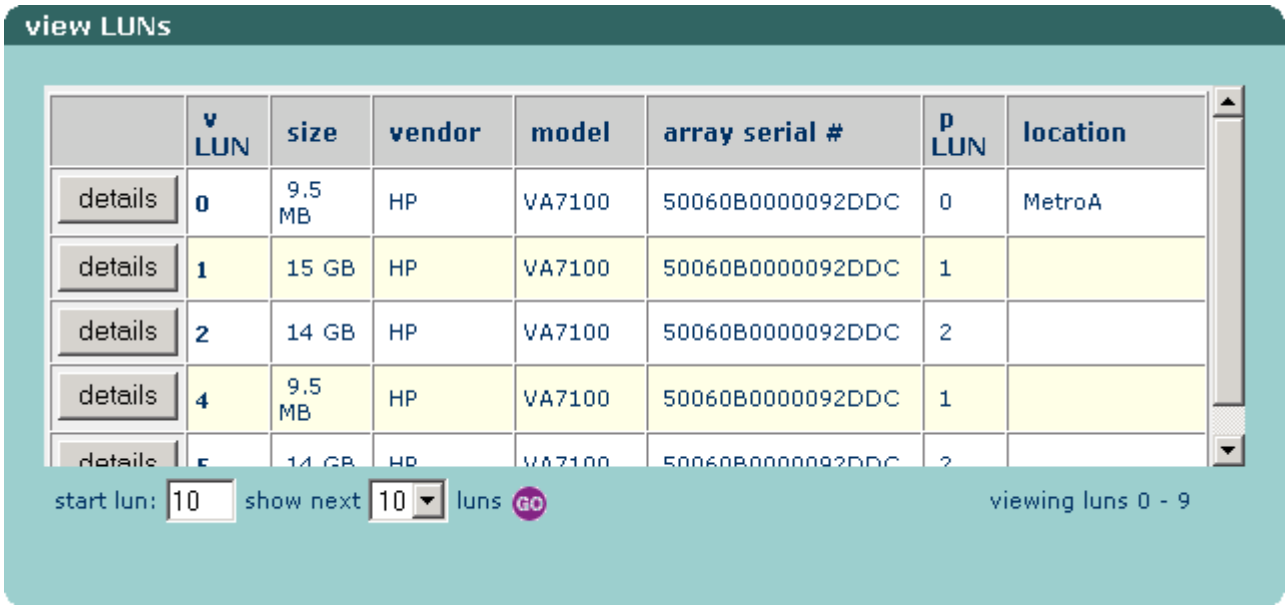

Figure 48: View LUNs window (with location)

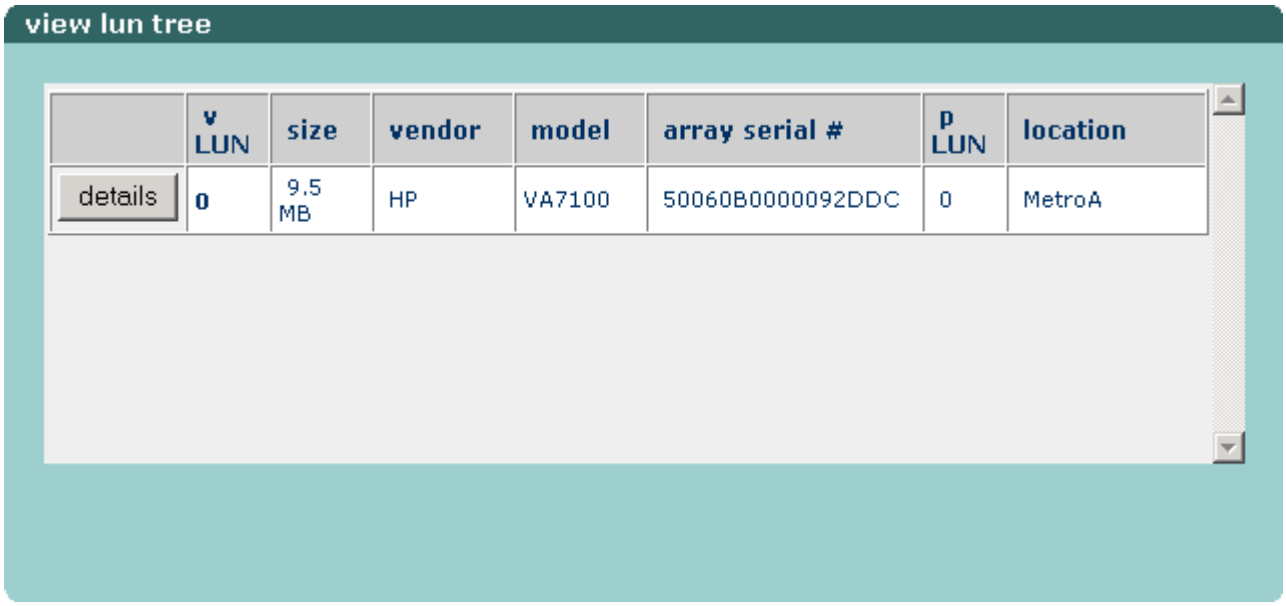

Figure 49: View lun tree window (with location)

Viewing LUN Statistics

You can view LUN statistics for read, write, and mirror requests. The statistics are measured by:

- **MB**—The size of the request in megabytes
- **I/O**—The number of input and output requests

To view LUN statistics:

1. Choose **LUNs > lun stats**.

   The **view luns** window opens, displaying statistics (Figure 50).
2. Select **info** for a LUN.

   The **lun x** window opens, displaying statistics (Figure 51).

The fields on the lun x window are:

- **mb read**—The amount of megabytes read
- **mb written**—The amount of megabytes written
- **i/o read**—The number of input and output requests for read requests
- **i/o written**—The number of input and output requests for write requests
- **mb mirrored**—The amount of megabytes mirrored
- **i/o mirrored**—The number of input and output requests for mirror requests

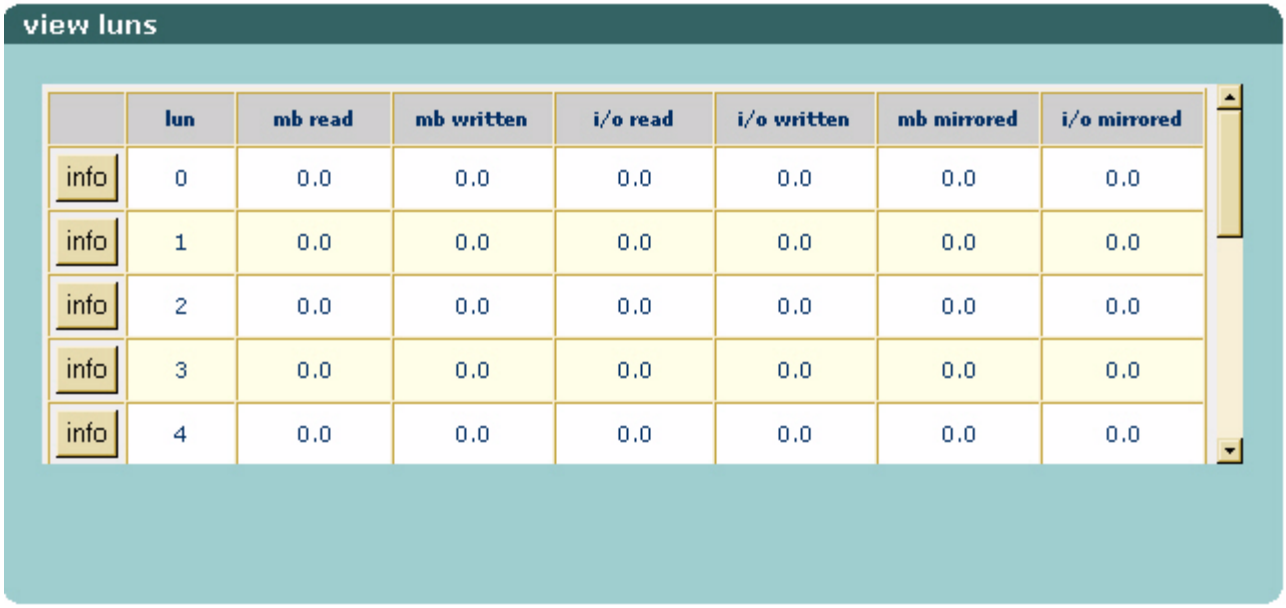

Figure 50: View luns window (displaying LUN statistics)

Figure 51: lun x window (displaying LUN statistics)

Viewing the LUN Matrix

You can determine how you are using LUNs by viewing the LUN matrix. The LUN matrix is an alternative to viewing the **lun x** window (Figure 39) and contains the following column headings:

- **partition**
- **partitioned**
- **pt image source**
- **pt image target**
- **expanded**
- **mapped**
- **fcp mirror source**
- **fcp mirror target**
- **ip mirror source**
- **ip mirror target**
- **hp service guard**
- **exported lun**
- **imported lun**

Columns marked with an "**x**" indicate how the LUN is being used. In Figure 52, you can see that LUN 9 has been partitioned and that LUN 5 and LUN 6 are mapped.

To view the LUN matrix:

Choose **LUNs > lun matrix**.

The **view luns** window opens, displaying the matrix (Figure 52).

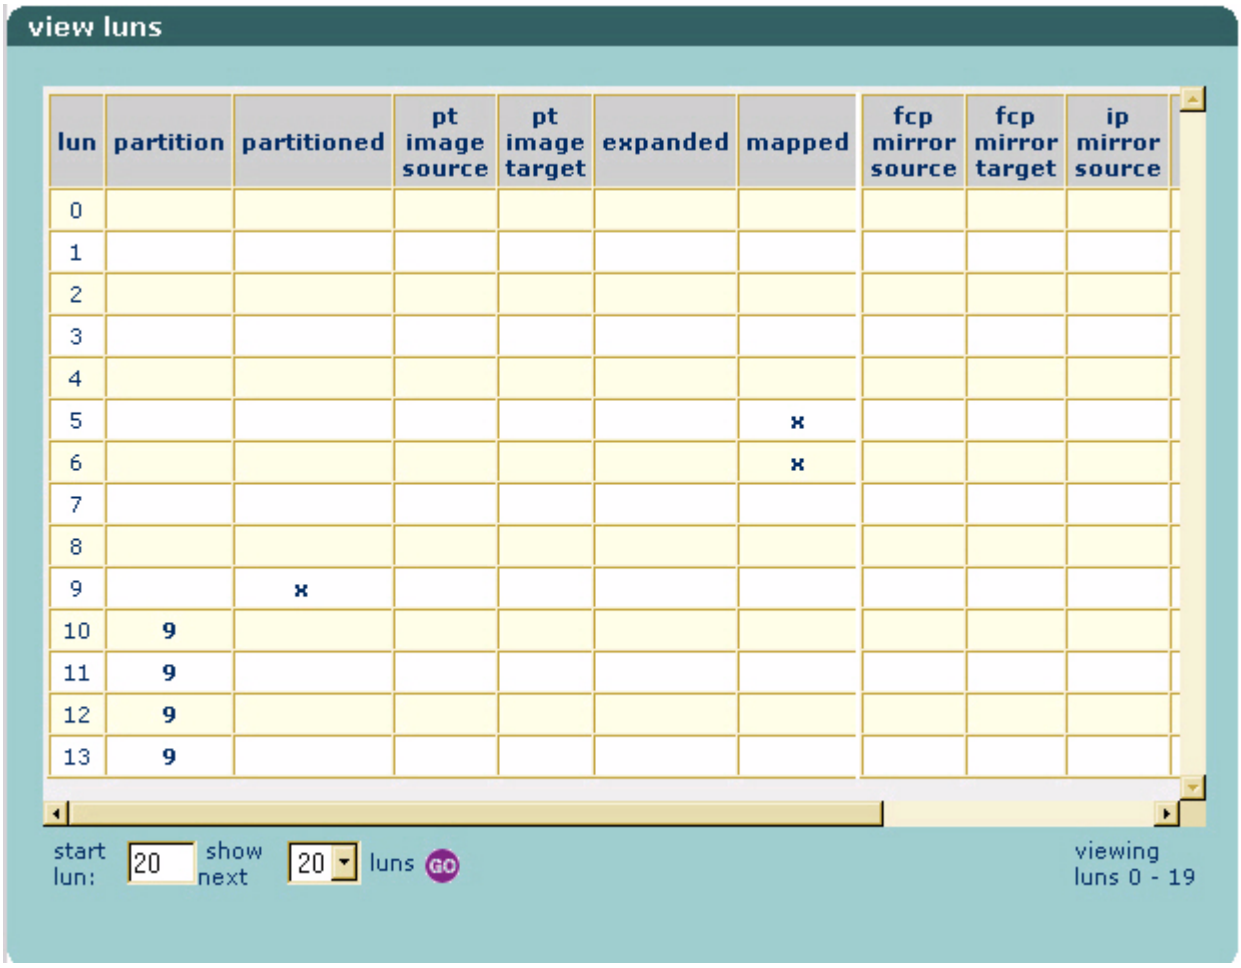

Figure 52: View luns window (displaying the LUN matrix)

Deleting LUNs

You may need to delete a physical LUN because:

- You have disconnected the LUN from the appliance.
- You want to increase the size of the LUN, which is a hardware LUN expansion. (See "LUN Expansion Types" on page 74 for more information.)

To delete a LUN:

1. Disconnect the LUN from the appliance.
2. Choose **LUNs > delete lun**.

   The **delete luns** window opens (Figure 53).
3. Choose the LUN you want to delete.
4. Click **delete**.

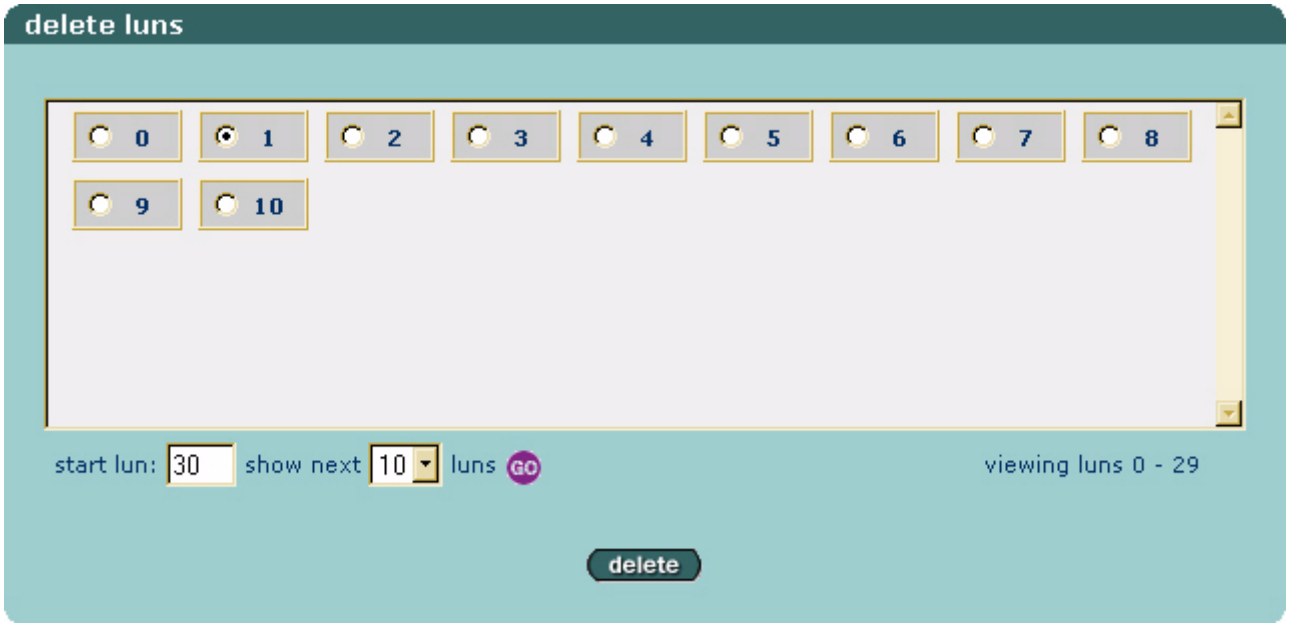

Figure 53: Delete luns window

Rescanning LUNs

A LUN rescan updates the appliance configuration with the most current LUN information.

You should perform a LUN rescan to when:

- You add or delete LUNs.
- You make changes to existing LUNs, such as creating partitions or expansions.

Note: If you add or delete several LUNs or make changes to several LUNs, the LUN rescan may take more time to complete. HP recommends that you schedule LUN rescans during off-peak hours or periods of low usage.

To perform a LUN rescan:

1. Choose **LUNs > rescan luns**.

   The LUN rescan begins.

2. When the rescan has completed, the **status** window opens, indicating the number of new LUNs detected (Figure 54).

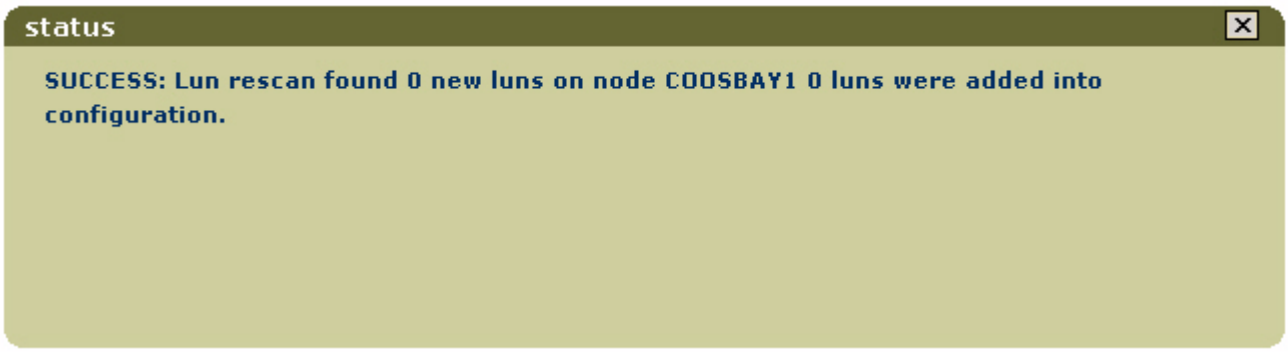

Figure 54: Status window

Managing User Access

8

This chapter describes how to manage the users who have access to the appliance. It covers the following topics:

- About Roles and Privileges, page 98
- Creating Roles, page 99
- Modifying Roles, page 101
- Deleting Roles, page 103
- Creating Users, page 104
- Changing User Passwords, page 106
- Deleting Users, page 107

About Roles and Privileges

Before you create a user, determine which privileges the user will have on the appliance. The appliance software provides over 200 privileges. Each privilege reflects an action that a user can perform using the user interface or the CLI. Once you determine a user's privileges, you can create a role comprised of those privileges. For example, you want to assign one user to manage LUN partitions and expansions. To create a role for that user, add all privileges for LUN partitions and expansions. Refer to "User Commands" in the *HP OpenView Continuous Access Storage Appliance Command Line Reference Guide* for a description of each privilege.

The appliance software includes two predefined roles:

- **Admin**—All privileges
- **Guest**—Viewing privileges only

Note: The appliance software creates the Guest role when HP service personnel execute the `initservice` command during appliance installation.

Creating Roles

To create a role:

1. Choose **utilities > roles**.

   The **role administration** window opens (Figure 55).

2. Enter a name for the role in the **create new role** field and click **create**.

   The **create role** window opens, displaying the privileges you can assign to the role (Figure 56).

3. To add privileges to a role:

   a. Select the privileges in the left column of the **create role** window.

   b. Click the right arrow to move the privileges to the right column.

4. To remove privileges from a role:

   a. Select the privileges in the right column of the **create role** window.

   b. Click the left arrow to move the privileges to the left column.

5. Click **save**.

   The new role is added to the list in the **existing roles** field (Figure 57).

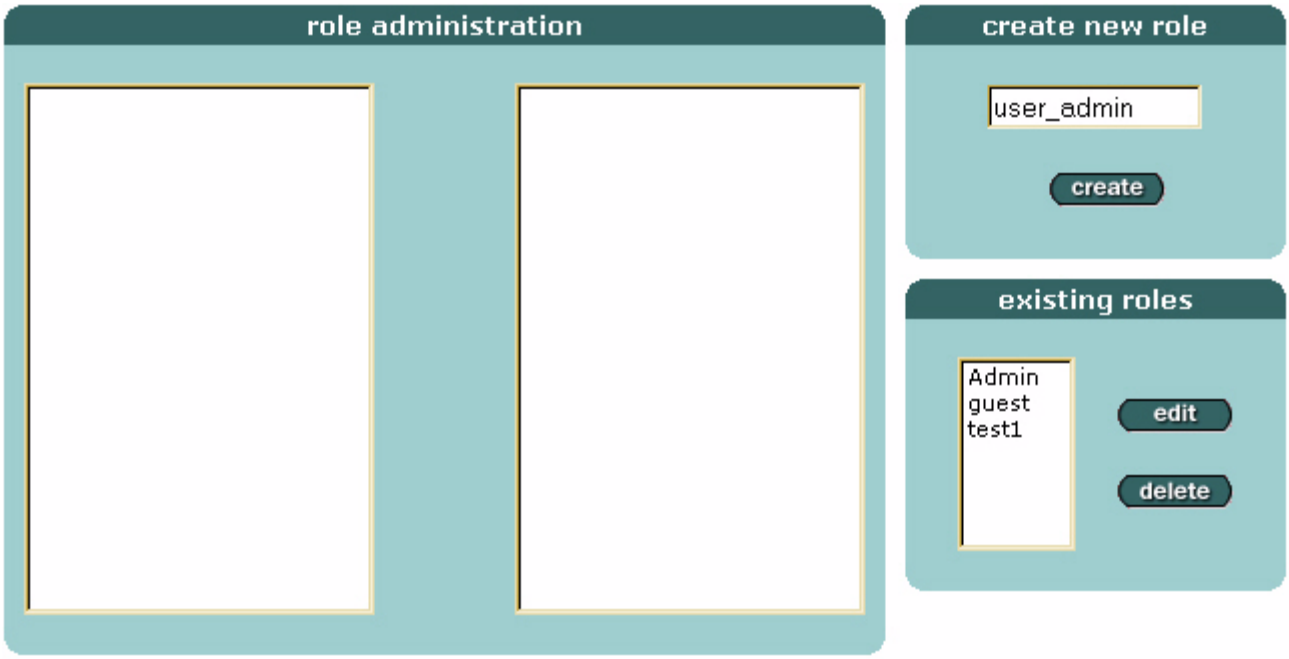

Figure 55: Role administration window

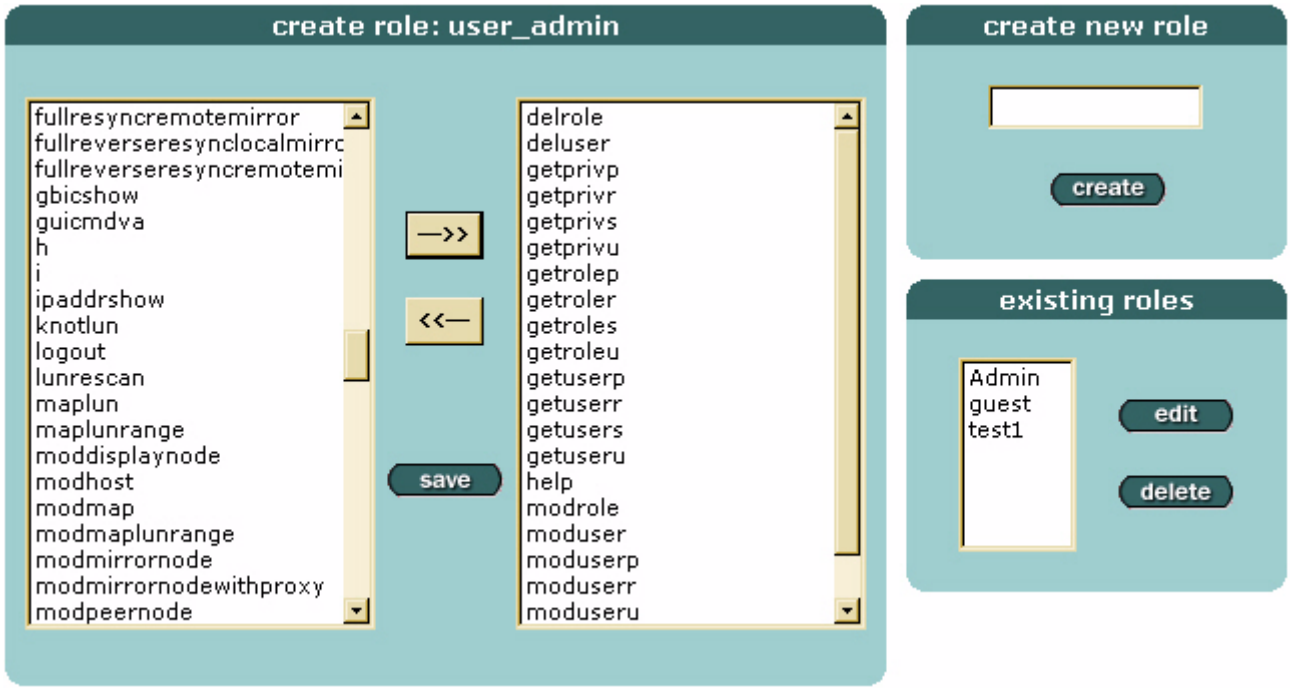

Figure 56: Create role window

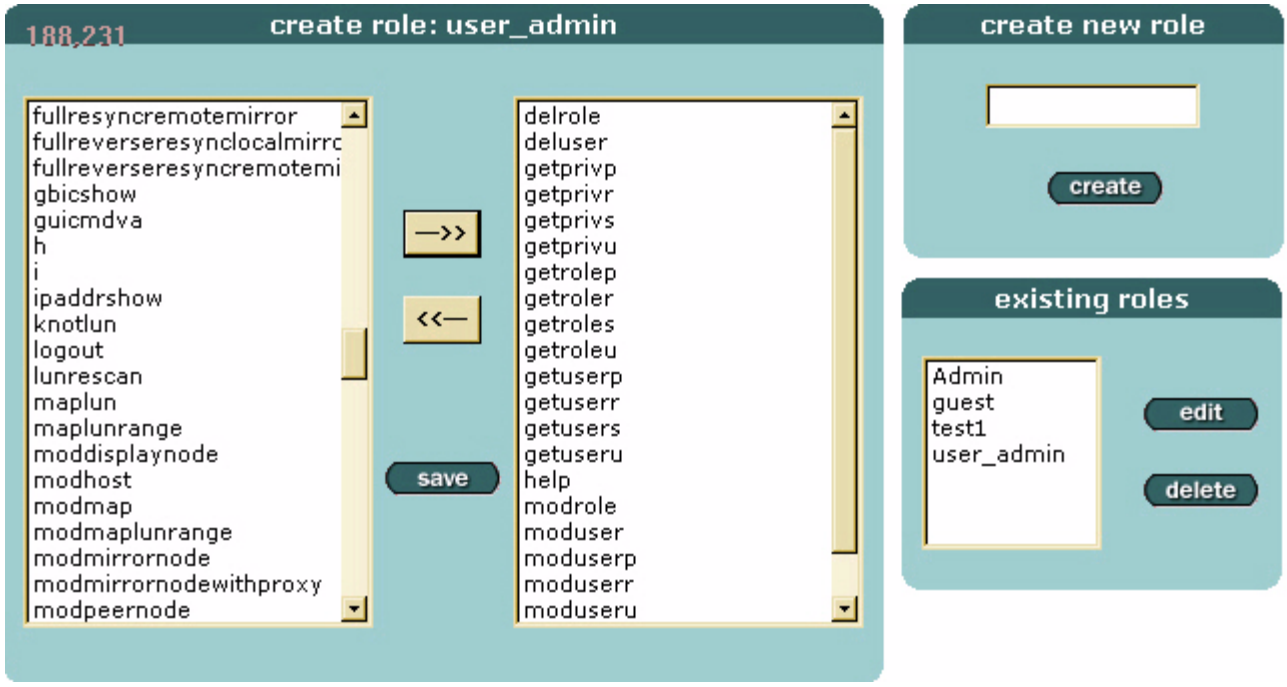

Figure 57: Create role window (new role added to existing roles)

Modifying Roles

You can modify the roles you assign to users. This section covers the following topics:

- Modifying User Roles
- Editing Role Privileges

Modifying User Roles

To modify a user's role:

1. Choose **utilities**.

   The **view users** window opens (Figure 62).
2. Click **modify** for the user whose role you want to change.

   The **user** window opens (Figure 63).
3. Select a role in the **role** field.
4. Click **save**.

Editing Role Privileges

To edit role privileges:

1. Choose **utilities > roles**.

   The **role administration** window opens (Figure 55).
2. Select the role you want to edit in the **existing roles** field and click **edit**.

   The **edit role** window opens (Figure 58).
3. To add privileges to a role:
   a. Select the privileges in the left column of the **edit role** window.
   b. Click the right arrow to move the privileges to the right column.
4. To remove privileges from a role:
   a. Select the privileges in the right column of the **edit role** window.
   b. Click the left arrow to move the privileges to the left column.
5. Click **save**.

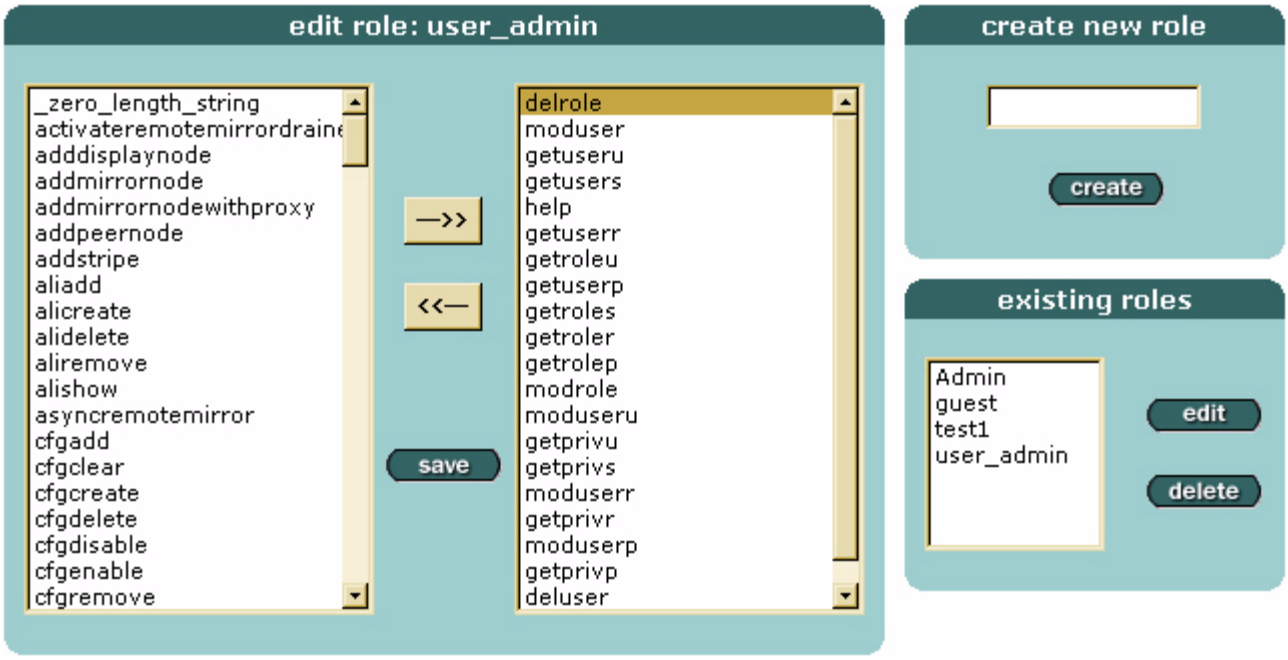

Figure 58: Edit role window

Deleting Roles

You can choose to delete a role that is no longer required. For example, you determine that a specific role is not used often, so you incorporate the privileges of that role into another role.

To delete a role:

1. Choose **utilities > roles**.

   The **role administration** window opens (Figure 55).

2. Select the role you want to delete in the **existing roles** field and click **delete**.

   The **delete role** window opens (Figure 59).

3. Click **delete**.

   The **role administration** window opens. The role you deleted is not listed in the **existing roles** field.

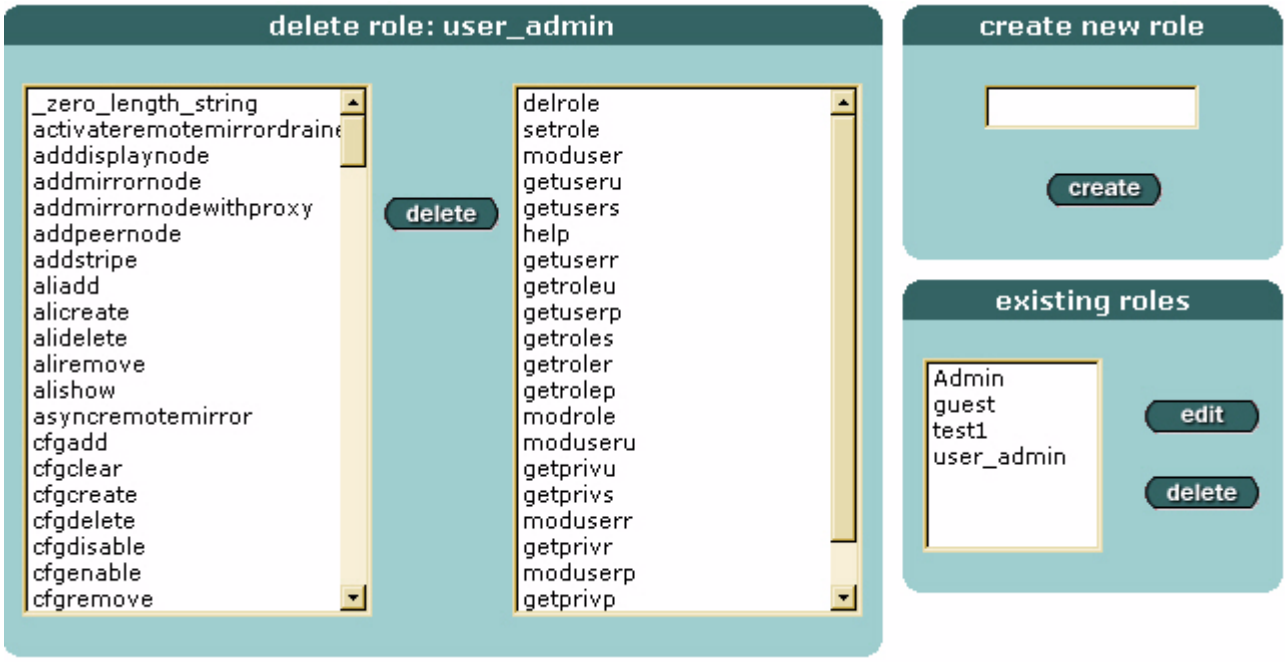

Figure 59: Delete role window

Creating Users

To create a user:

1. Choose **utilities**.

   The **view users** window opens (Figure 60).

2. Click **create new user**.

   The **create new user** window opens (Figure 61).

3. Enter a name for the user in the **username** field.

4. Enter a password in the **password** field.

5. Reenter the password in the **confirm** field.

   If you do not assign a password to the user or you create a password that begins with *nopass*, the user cannot access the CLI from the user interface. However, the user can still access the CLI from the CASA console or another workstation.

Note: HP recommends that you assign a unique password to each user.

6. Select a role for the user in the **role** field. (See "About Roles and Privileges" on page 98 for more information.)

7. Click **save**.

   The **view users** window opens, displaying the new user (Figure 62).

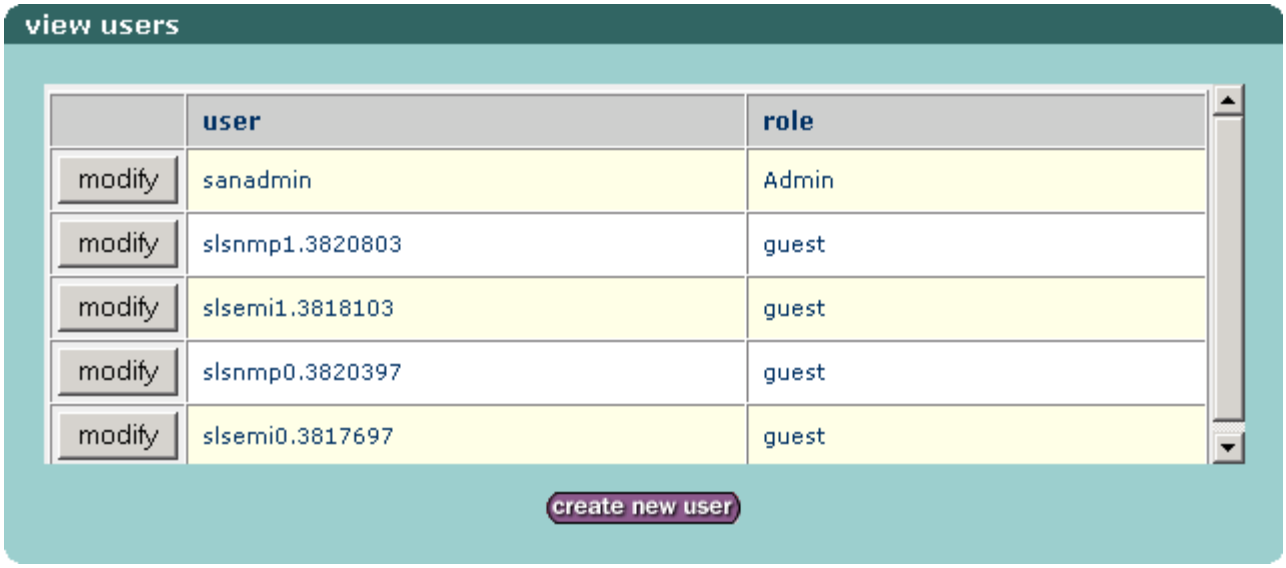

Figure 60: View users window

The **host** field on the **create new user** window shows the management IP address of the local node. You use this IP address to manage the appliance through the user interface and the CLI. When you create a new user on one node, the appliance software adds that user to the peer node.

Figure 61: Create new user window

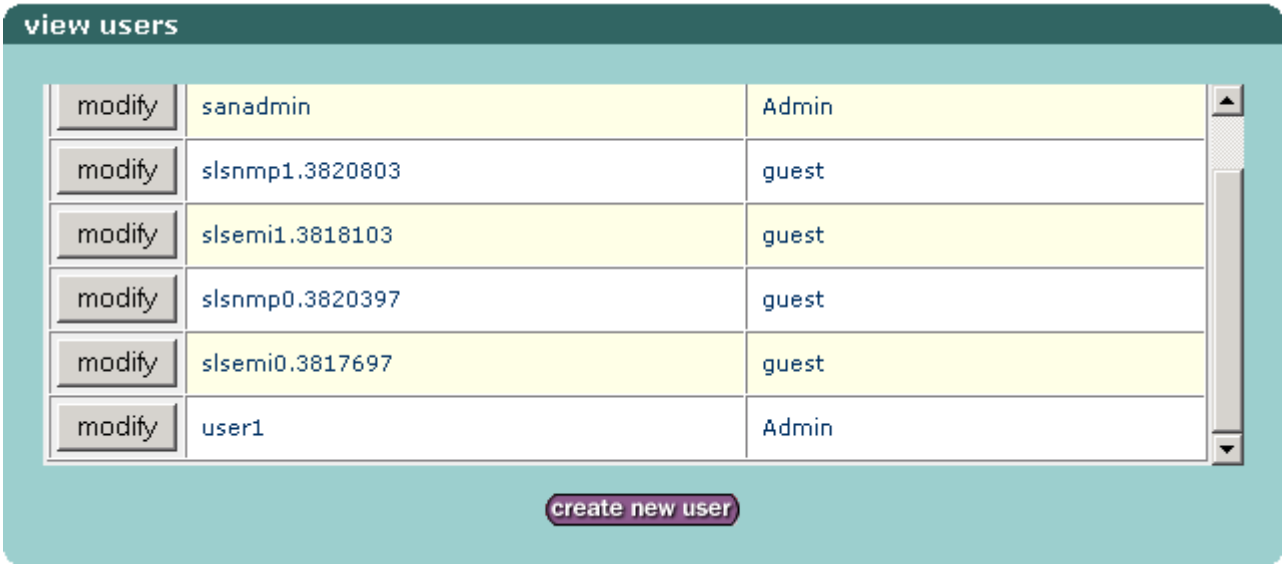

Figure 62: View users window (displaying the new user)

Changing User Passwords

To change a user's password:

1. Choose **utilities**.

   The **view users** window opens (Figure 62).

2. Click **modify** for the user whose password you want to change.

   The **user** window opens (Figure 63).

3. Click **(change password)**.

   The **user** window reopens, displaying the **current password**, **new password**, and **confirm password** fields (Figure 64).

4. Enter the user's current password in the **current password** field.
5. Enter the new password in the **new password** and **confirm password** fields.
6. Click **change password**.

Figure 63: User window

Figure 64: User window (displaying the password fields)

Deleting Users

For security purposes, HP recommends that you delete users who no longer require access to the appliance.

To delete a user:

1. Choose **utilities > delete users**.

   The **view users** window opens (Figure 65).

2. Choose the user you want to delete.
3. Click **delete**.

   The **view users** window reopens. The user you deleted is not listed.

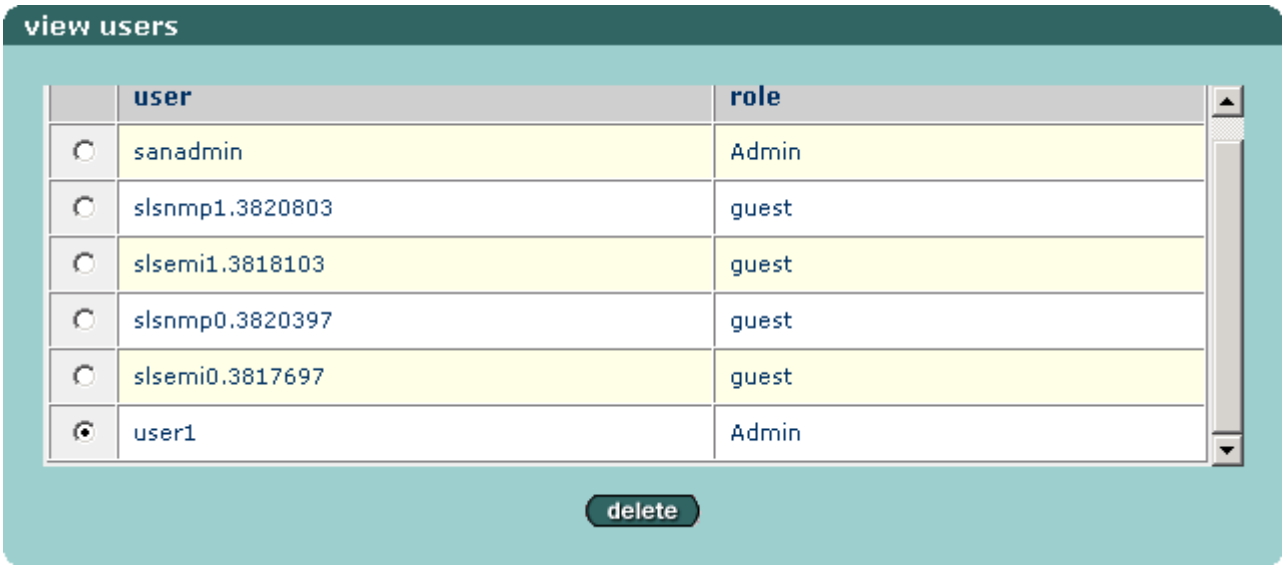

Figure 65: View users window (deleting a user)

Managing Local FCP Mirrors

9

This chapter defines local Fibre Channel Protocol (FCP) mirroring and describes how to implement it in your network environment. It covers the following topics:

- About Local FCP Mirrors, page 110
- Local FCP Mirror Details, page 111
- Creating a Local FCP Mirror, page 112
- About Mirror States, page 114
- Pausing Local FCP Mirrors, page 116
- Resynchronizing Local FCP Mirrors, page 118
- Deleting Local FCP Mirrors, page 120

About Local FCP Mirrors

Local FCP mirroring is data replication between two storage arrays connected to the same appliance through a Fibre Channel Protocol (FCP) link. The FCP link enables data replication between storage arrays located up to 10 kilometers apart. For example, you can mirror data between storage arrays located on different floors of the same building.

Local FCP mirroring enables you to maintain a complete, identical image of data on two separate disks. If an error disables either storage array, the appliance automatically forwards network traffic to the operational storage array without interruption. After you repair or replace the disabled storage array, the appliance automatically copies data from the working image to the image on the newly operational storage array.

Note: The target LUN of a local FCP mirror must be the same size or greater than the source LUN.

Local FCP Mirror Details

To create a local FCP mirror, you must specify:

- **Source LUN**—The storage disk containing the data you want to replicate. See "Rules for Source and Target LUNs" for more information.
- **Target LUN**—The storage disk containing the replicated data. See "Rules for Source and Target LUNs" for more information.
- **State**—The state of the source LUN and target LUN at the time you create the mirror. Choose **full resync** to ensure the data on both LUNs is the same before replication begins. Choose **no resync** to enable replication immediately; note that existing data on the source LUN and target LUN may not be the same.

The appliance performs FCP mirroring in synchronous mode—before the appliance software marks a replication request successful, the source and target LUNs must return a write-success message.

Note: You can mirror data from one source LUN to a maximum of nine target LUNs.

Rules for Source and Target LUNs

When you choose the LUNs for your source and target, ensure that:

- You do not use the same LUN for the source and the target.
- The source LUN is not a target LUN for another mirror.
- The target LUN is not a source LUN or a target LUN for another mirror.
- The target LUN is not a mapped LUN. See "Managing LUNs" on page 61 for more information about mapped LUNs.
- The LUN you choose for either the source or the target is not a point-in-time image target. See "Managing Point-in-Time Images" on page 185 for information about point-in-time image targets.
- The LUN you choose for either the source or the target is not an exported LUN (either *exported to* or *exported from* another appliance). See "Managing Cross-Appliance FCP Mirrors" on page 121 for more information about exported LUNs.

Creating a Local FCP Mirror

To create a local FCP mirror:

1. Choose **mirrors > fcp mirror > create fcp mirror**.

   The **create fcp mirror** window opens (Figure 66).
2. Choose the source LUN in the **select source lun** field.
3. Choose the target LUN in the **select target lun** field.
4. Choose **no resync** or **full resync** in the **select state** field.
5. Click **submit**.

   The **view fcp mirrors** window opens (Figure 67). If you chose **no resync**, the mirror state is Active. If you chose **full resync**, the mirror state is Resynchronizing.

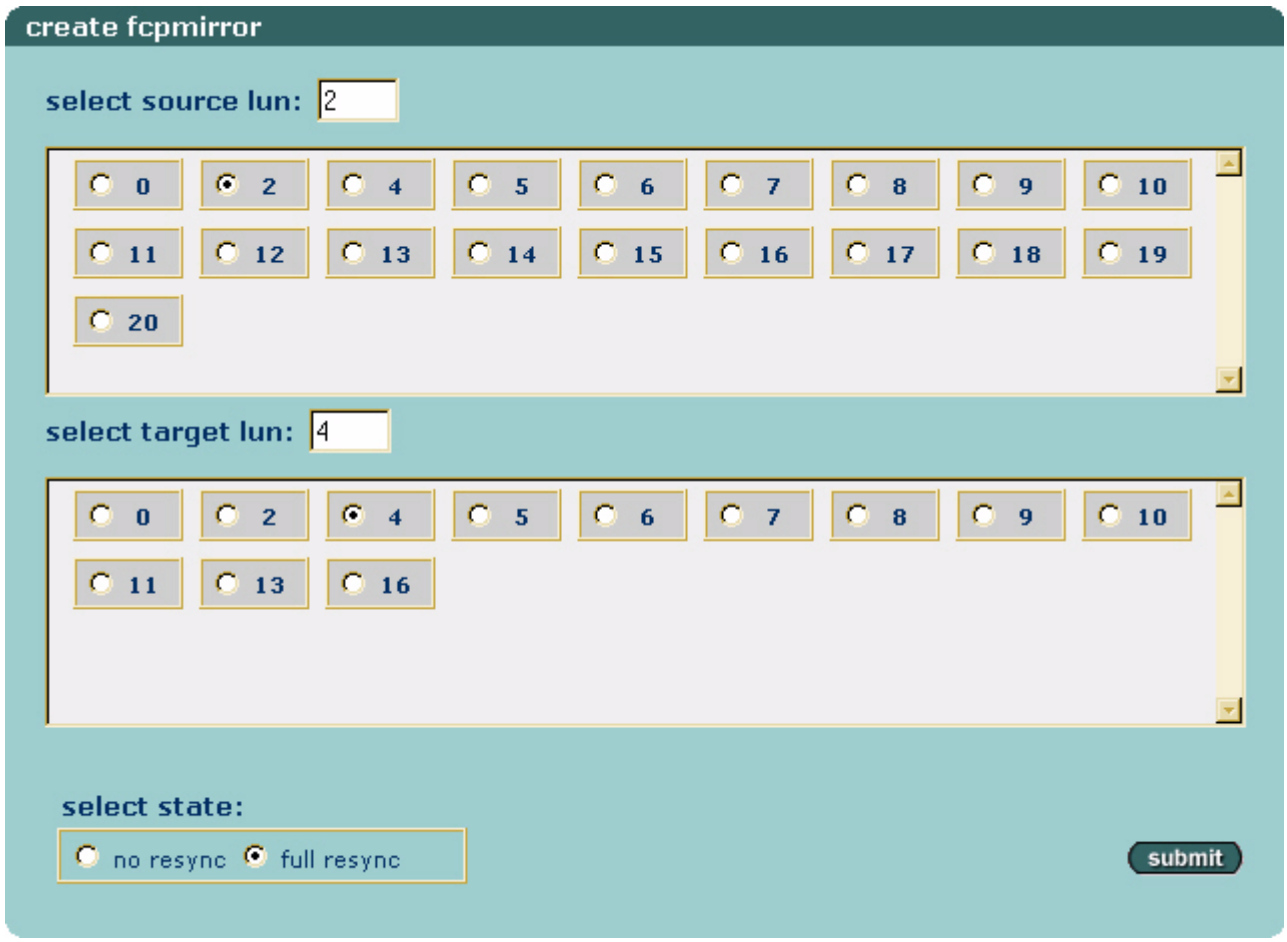

Figure 66: Create fcp mirror window

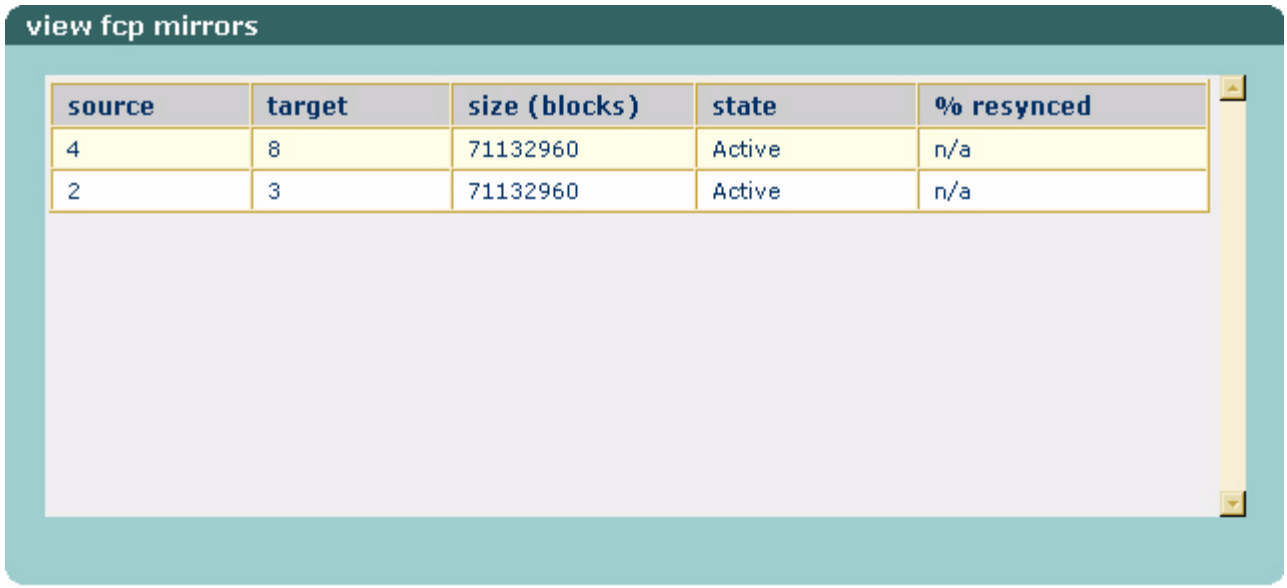

Figure 67: View fcp mirrors window

Table 5 describes each field in the **view fcp mirrors** window.

Table 5: Fields in the view fcp mirrors window

Field	Description
source	The source LUN of the mirror.
target	The target LUN of the mirror.
size (blocks)	The size (in blocks) of the mirrored LUNs.
state	The current state of the mirror. See "About Mirror States" on page 114 for more information.
% resynced	The percentage of the journal file that has been processed during a resynchronization. Refresh the window periodically to update the resynchronization's progress. The field value is 100% when either of the following occurs: ■ The resynchronization is complete. ■ You create the mirror using the **no resync** option.

About Mirror States

You can monitor local FCP mirrors by checking the **state** field in the **view fcp mirrors** window.

Table 6 describes the mirror states for FCP mirrors.

Table 6: Mirror states for FCP mirrors

Mirror State	Description
Active	The state of normal FCP mirroring. The appliance software processes replication requests on the source LUN and sends them to the target LUN for replication. From this state, the mirror can enter the Inactive state or Paused state.
Full Down	Both the source and target LUNs are down.
Inactive	The mirror becomes inactive if the source or target LUN goes down.
Paused	Mirroring stops. The appliance software tracks all new replication requests in the mirror's journal file.
Resynchronizing	When the mirror is resynchronizing, one of the two appliance nodes copies data from one LUN to the other LUN for that mirror. I/O sent to the node in the Resynchronizing state for that mirror is handled locally. When resynchronization is complete, the mirror enters the Active state. If you chose **full resync** when creating the mirror, the mirror automatically enters the Resynchronizing state.
Waiting	When the mirror is resynchronizing, the node that is not copying data (see Resynchronizing) enters the Waiting state. If I/O is sent to the node in the Waiting state, the appliance software forwards the I/O to the node performing the resynchronization.

Figure 68 illustrates the transitions between FCP mirror states.

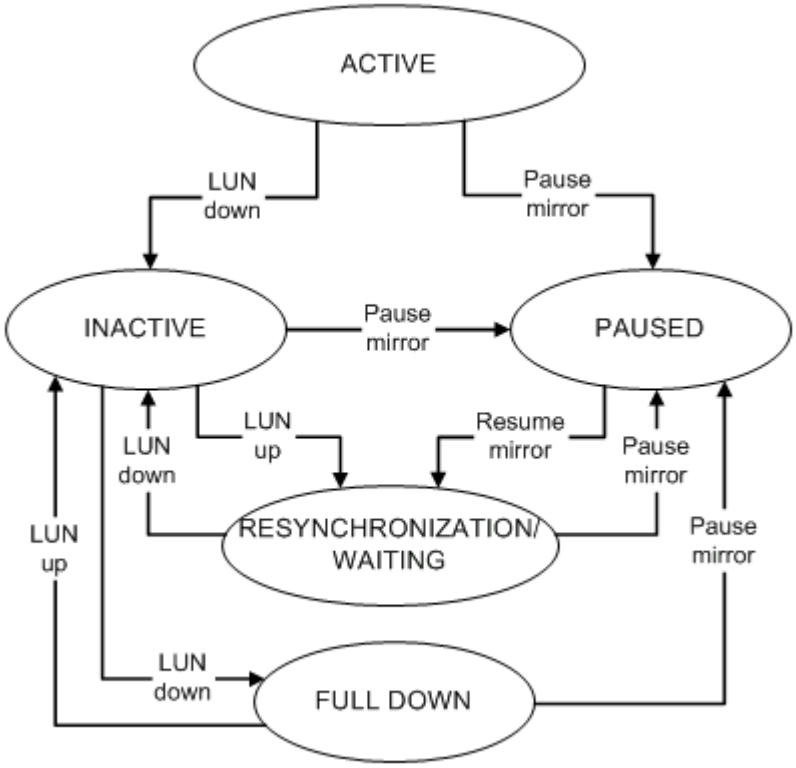

Figure 68: FCP mirror states and transitions

Pausing Local FCP Mirrors

Pausing a local FCP mirror temporarily stops data replication between the source LUN and the target LUN. While the mirror is paused, the appliance software tracks new replication requests in the mirror's journal file.

Note: Disk redundancy is disabled for paused mirrors.

You can pause a mirror to save an image of the data on the target LUN. After you back up the data, resume mirroring.

To pause a local FCP mirror:

1. Choose **mirrors > fcp mirror > pause-resync**.

   The **pause resync fcp mirrors** window opens (Figure 69).

2. Choose the mirror you want to pause.

   The **selected mirror** window opens (Figure 70).

3. Choose **pause** and click **submit**.

   The **pause resync fcp mirrors** window opens, showing the state is Paused (Figure 71).

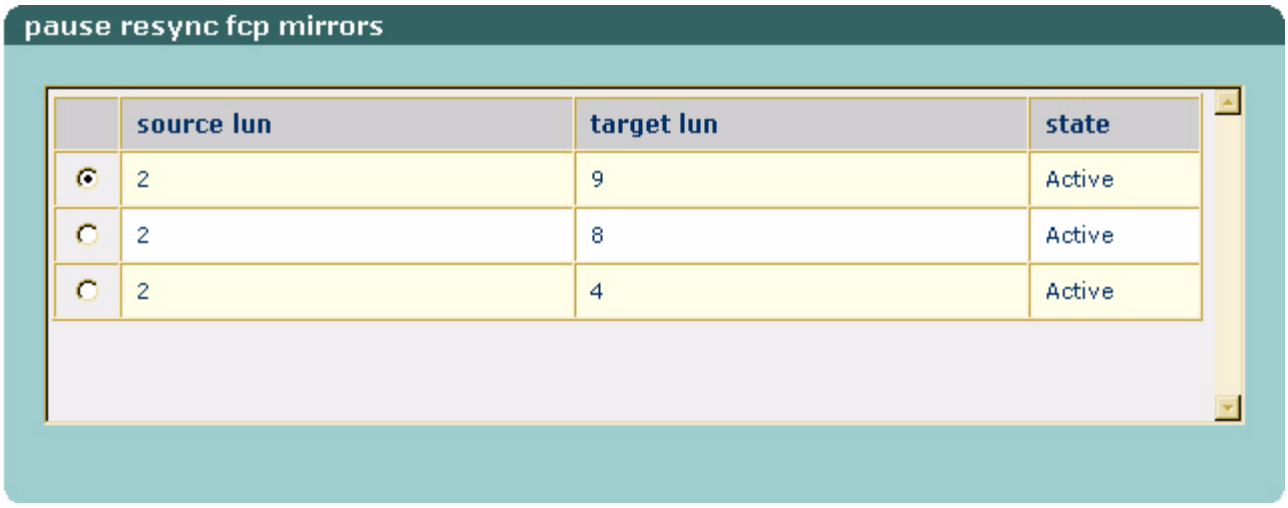

Figure 69: Pause resync fcp mirrors window

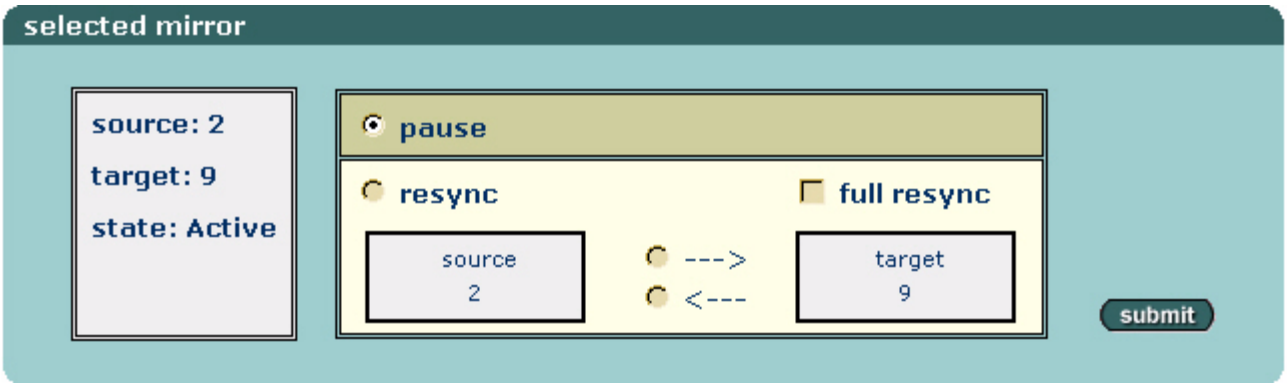

Figure 70: Selected mirror window

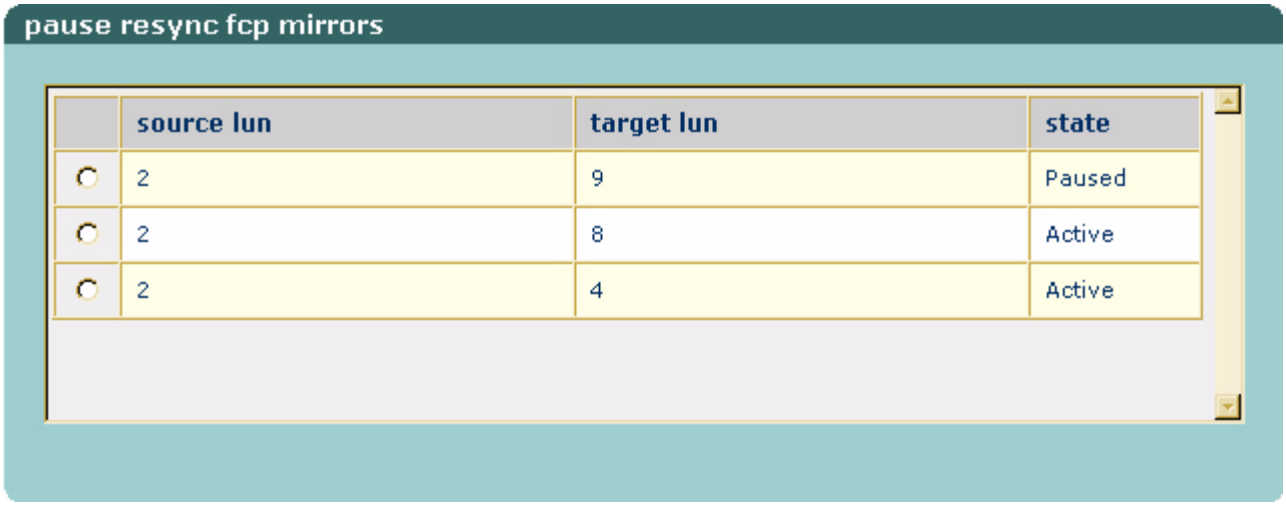

Figure 71: Pause resync fcp mirrors window (Paused mirror state)

Resynchronizing Local FCP Mirrors

Mirroring can stop when:

- A failure occurs on either the source LUN or the target LUN. (See "About Mirror States" on page 114.)
- You pause a mirror. (See "Pausing Local FCP Mirrors" on page 116.)

Before you resume mirroring, you must resynchronize the data on the source and target LUNs so that both LUNs contain the same data. You must choose:

- **Data to be resynchronized**—Choose **resync** to only copy the data that changed since you paused the mirror. Choose **full resync** to copy all of the data (the entire LUN).
- **Direction of the resynchronization**—Choose forward resynchronization (source to target) to update the target LUN with the source LUN's data. Choose reverse resynchronization (target to source) to update the source LUN with the target LUN's data. In either direction, the source LUN does the work—it sends the data (forward resynchronization) or obtains the data (reverse resynchronization).

Note: When you choose the resynchronization direction, be sure that you know which LUN has the most current data. Once you choose a direction, the data on the LUN you are updating will be overwritten.

To resynchronize a local FCP mirror:

1. Choose **mirrors > fcp mirror > pause-resync**.

   The **pause resync fcp mirrors** window opens.
2. Choose the mirror you want to resynchronize.

   The **selected mirror** window opens (Figure 72).
3. Choose **resync** or **full resync**.
4. Choose the resynchronization direction–from source to target (forward) or from target to source (reverse).
5. Click **submit**.

   The **pause resync fcp mirrors** window opens. The mirror you resynchronized shows the mirror state as Resynchronizing (Figure 73).

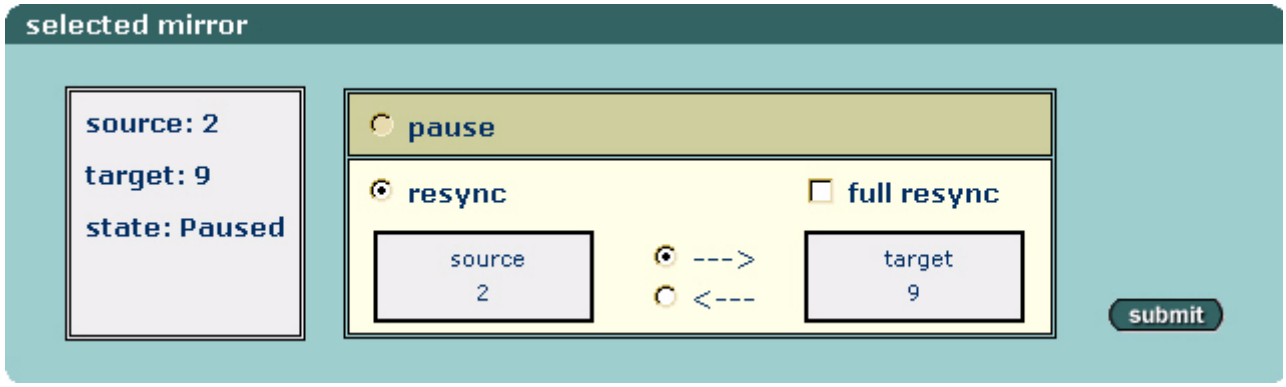

Figure 72: Selected mirror window

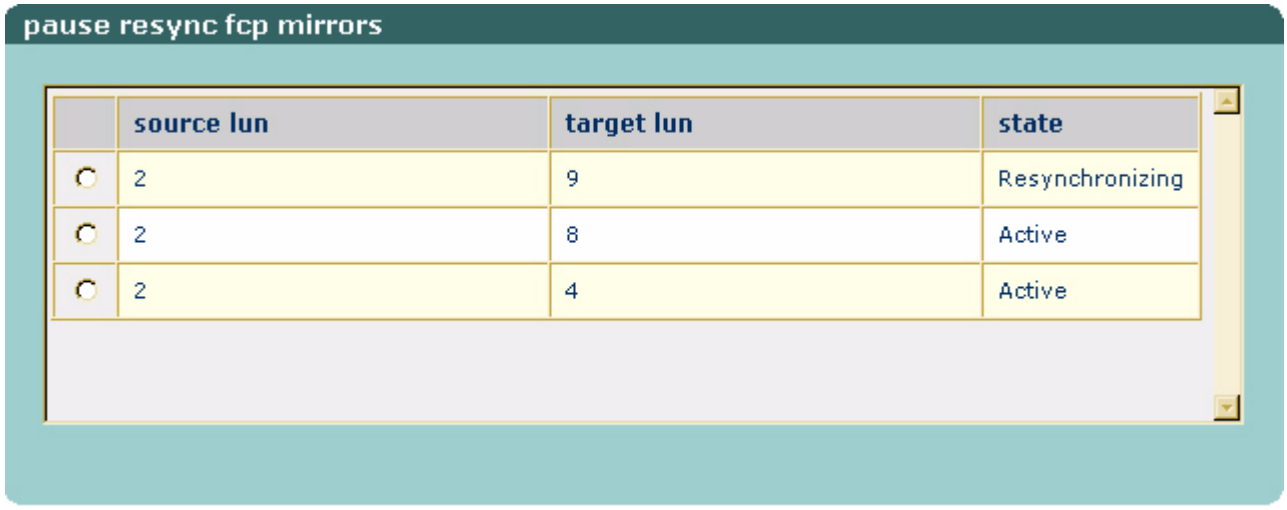

Figure 73: Pause resync fcp mirrors window (Resynchronizing mirror state)

When the mirror has finished resynchronizing, the mirror state returns to Active. You can check the progress of the resynchronization on the **view fcp mirrors** window (Figure 74). The **% resynced** field indicates how much of the resynchronization is complete.

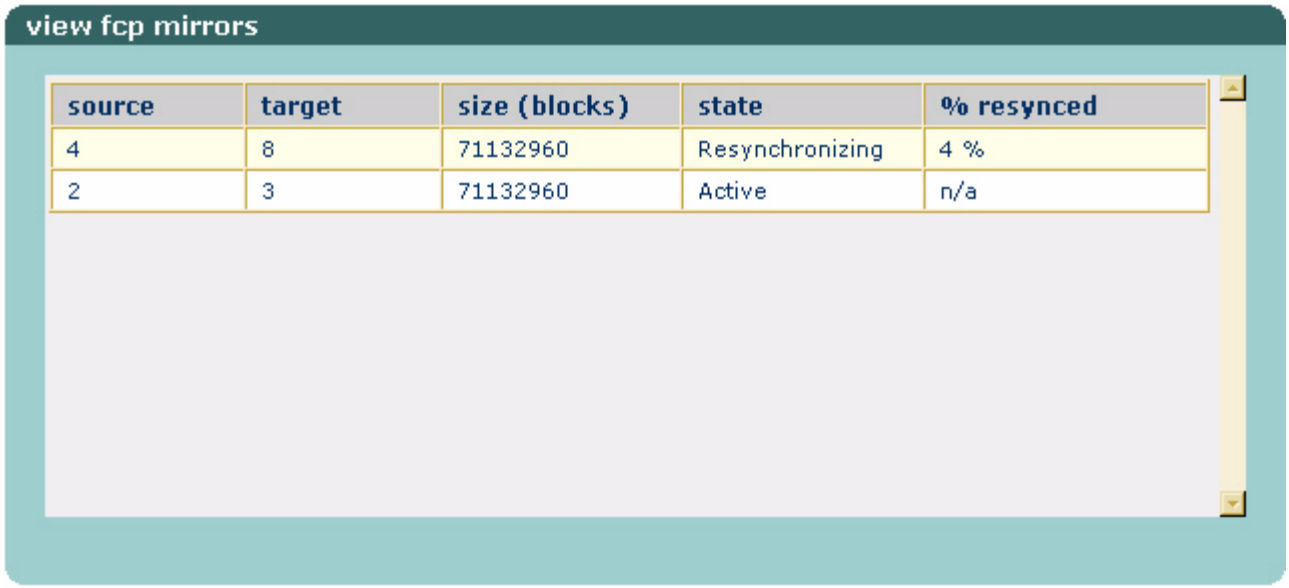

Figure 74: View fcp mirrors (Resynchronizing mirror state)

Deleting Local FCP Mirrors

You delete a local FCP mirror when you no longer need to replicate data.

To delete a local FCP mirror:

1. Choose **mirrors > fcp mirror > delete fcp mirror**.

   The **delete fcp mirrors** window opens (Figure).
2. Choose the FCP mirror you want to delete.
3. Click **submit**.

   The **view fcp mirrors** window opens. The mirror you deleted is not listed.

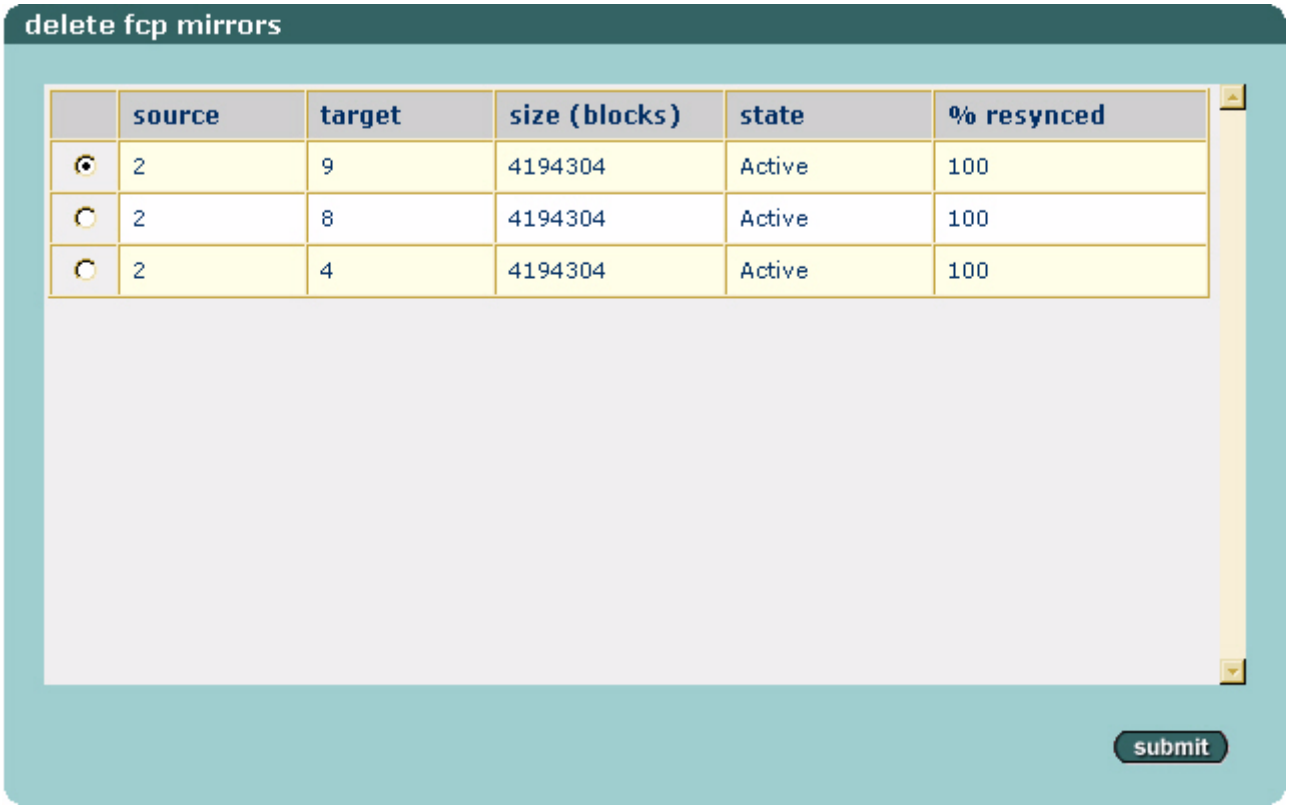

Figure 75: Delete fcp mirrors window

Managing Cross-Appliance FCP Mirrors

10

This chapter defines cross-appliance FCP mirroring and describes how to implement it in your network environment. It covers the following topics:

- About Cross-Appliance FCP Mirrors, page 122
- Creating a Cascaded Configuration, page 123
- Exporting LUNs, page 126
- Creating Cross-Appliance FCP Mirrors, page 137
- About Mirror States, page 140
- Pausing Cross-Appliance FCP Mirrors, page 141
- Resynchronizing Cross-Appliance FCP Mirrors, page 143
- Deleting Cross-Appliance FCP Mirrors, page 145

About Cross-Appliance FCP Mirrors

Cross-appliance FCP mirroring is data replication between appliances that are connected through a cascaded Fibre Channel Protocol (FCP) link and are separated by long distances. For example, you can mirror data between cascaded appliances that reside in separate racks within the same room, in different rooms at the same location, or in two separate locations at distances supported by FC networks.

Note: The Continuous Access Storage Appliance supports a maximum of two cascaded appliances.

Cross-appliance FCP mirroring enables you to maintain a complete, identical image of data on disks in two locations. If an error disables one of the appliances or a site failure occurs, you can forward network traffic to the operational appliance and storage device without losing data access. After repairing or replacing the disabled appliance, copy the data from the working image to the image on the newly operational appliance.

Note: The target logical unit number (LUN) of a cross-appliance FCP mirror must be the same size or greater than the source LUN.

Creating a Cascaded Configuration

This section describes how to create a cascaded configuration. It covers the following topics:

- Wiring Cascaded Appliances
- Identifying Cascaded Nodes

Wiring Cascaded Appliances

You can create a cascaded configuration using switches and fabrics or connecting the appliances directly. Use dedicated ports as follows ((Figure 76):

- Fibre Channel connection from initiator port 9 (I5) on each node of Appliance A to target port 11 (T5) on each node of Appliance B (Node 0 to Node 0 and Node 1 to Node 1).
- Fibre Channel connection from initiator port 9 (I5) on each node of Appliance B and to target port 11 (T5) on each node of Appliance A (Node 0 to Node 0 and Node 1 to Node 1).

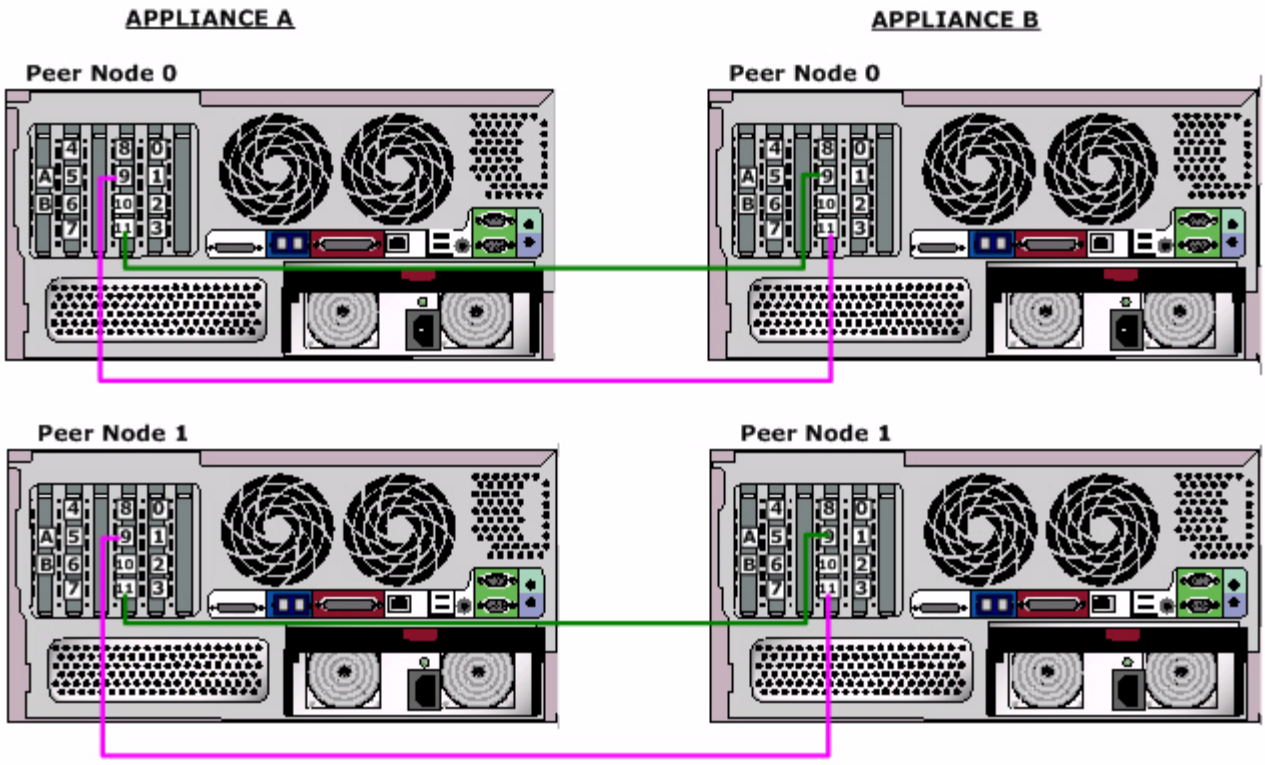

Figure 76: Port connections for cascaded appliances

Identifying Cascaded Nodes

After the appliances are cascaded, Node 0 of Appliance A identifies Node 0 of Appliance B as a host. In return, Node 0 of Appliance B identifies Node 0 of Appliance A as a host. This also applies to Node 1 on both appliances—each Node 1 identifies the other Node 1 as a host.

To ensure that each node properly identifies the other nodes:

1. Shut down the nodes on each appliance.
2. Connect the appliances as described in "Wiring Cascaded Appliances".
3. Boot Node 0 on Appliance A. Ensure that the node has initialized before continuing with step 4.

   To determine if the node has initialized, open sanosadmin (select **sanosadmin > launch** from the console). If the `Server [localhost]:` prompt is displayed, initialization is not complete. Close sanosadmin and wait a few minutes before launching sanosadmin again. If the `login:` prompt is displayed, initialization is complete.
4. Boot Node 0 on Appliance B. Ensure that the node has initialized before continuing with step 5.
5. From Node 0 on either appliance, open sanosadmin and enter the following command:

   ```
   get host
   ```

   Figure 77 shows the command output from Node 0 on Appliance A. Host 5 is Node 0 on Appliance B. The **OS** field displays `SAN.OS`, the appliance's operating system.

```
Hostid  Host Name         World Wide Name  Port Mask  Node ID  OS
______  ________________  _______________  _________  _______  _______________
1       210100E08B22492D  210100e08b22492d 0x0000007     0     UNKNOWN
4       210000E08B22B53C  210000e08b22b53c 0x0000007     1     UNKNOWN
5       COOSBAY2(5)       210000e08b05419a 0x0000007     0     SAN.OS_5.5.0.44
```

Figure 77: Output from Node 0 on Appliance A

6. Go to Node 0 on the other appliance (in this example, Appliance B) and repeat step 5.

   The output is similar to that shown in Figure 77 but is specific to Node 0 on Appliance B (Figure 78). Host 3 is Node 0 on Appliance A.

```
Hostid  Host Name        World Wide Name  Port Mask  Node ID  OS
______  ________________ _______________  _________  _______  _______________
1       210000E08B22AB60 210000e08b22ab60 0x0000007     1     UNKNOWN
3       COOSBAY1(5)      210000e08b05c8fe 0x0000007     0     SAN.OS_5.5.0.44
5       210000E08B221C2D 210000e08b221c2d 0x0000007     0     UNKNOWN
```

Figure 78: Output from Node 0 on Appliance B

7. Boot Node 1 on Appliance A. Ensure that the node has initialized before continuing with step 8.
8. Boot Node 1 on Appliance B. Ensure that the node has initialized before continuing with step 9.

9. From Node 1 on either appliance, enter the following command:

   ```
   get host
   ```

   Figure 79 shows the command output from Node 1 on Appliance A. Host 3 is Node 1 on Appliance B. Host 5 is Node 0 on Appliance B.

```
Hostid  Host Name        World Wide Name   Port Mask  Node ID OS
______  _______________  ________________  _________  _______ _______________
1       210100E08B22492D 210100e08b22492d 0x0000007    0      UNKNOWN
3       COOSBAY3(5)      210000e08b05439a 0x0000007    1      SAN.OS_5.5.0.44
4       210000E08B22B53C 210000e08b22b53c 0x0000007    1      UNKNOWN
5       COOSBAY2(5)      210000e08b05419a 0x0000007    0      SAN.OS_5.5.0.44
```

Figure 79: Output from Node 1 on Appliance A

10. Go to Node 1 on the other appliance (in this example, Appliance B) and repeat step 9.

    The output is similar to that shown in Figure 79 but is specific to Node 1 on Appliance B (Figure 80). Host 2 is Node 1 on Appliance A. Host 3 is Node 0 on Appliance A.

```
Hostid  Host Name        World Wide Name   Port Mask  Node ID OS
______  _______________  ________________  _________  _______ _______________
1       210000E08B22AB60 210000e08b22ab60 0x0000007   1      UNKNOWN
2       COOSBAY4(4)      210000E08b05409a 0x0000007   1      SAN.OS_5.5.0.44
3       COOSBAY1(5)      210000e08b05c8fe 0x0000007   0      SAN.OS_5.5.0.44
5       210000E08B221C2D 210000e08b221c2d 0x0000007   0      UNKNOWN
```

Figure 80: Output from Node 1 on Appliance B

If you follow these steps and the nodes cannot identify each other, either:

- You did not connect the nodes properly.
- You did not boot the nodes one at a time. For example, if you boot Node 0 on both appliances at the same time, the identification process may not succeed. In this case, you must perform a manual rescan.

Exporting LUNs

This section defines exported LUNs and describes how to export LUNs. It covers the following topics:

- About Exported LUNs
- Creating Exported LUNs
- Viewing Exported LUNs

About Exported LUNs

In a cascaded configuration, appliances cannot view the storage arrays connected to the other appliance. You must map the source and target LUNs between the appliances, which is called exporting LUNs (Figure 81). Exporting a LUN is the same as mapping a LUN to an appliance's local host. Mapping a LUN enables a host to access that LUN on a storage array. (See "Mapping LUNs" on page 63.) Mapping LUNs between appliances enables you to create a mirror between LUNs connected to different appliances.

Note: Before you create a cross-appliance mirror. you must export the appropriate LUNs between the cascaded appliances.

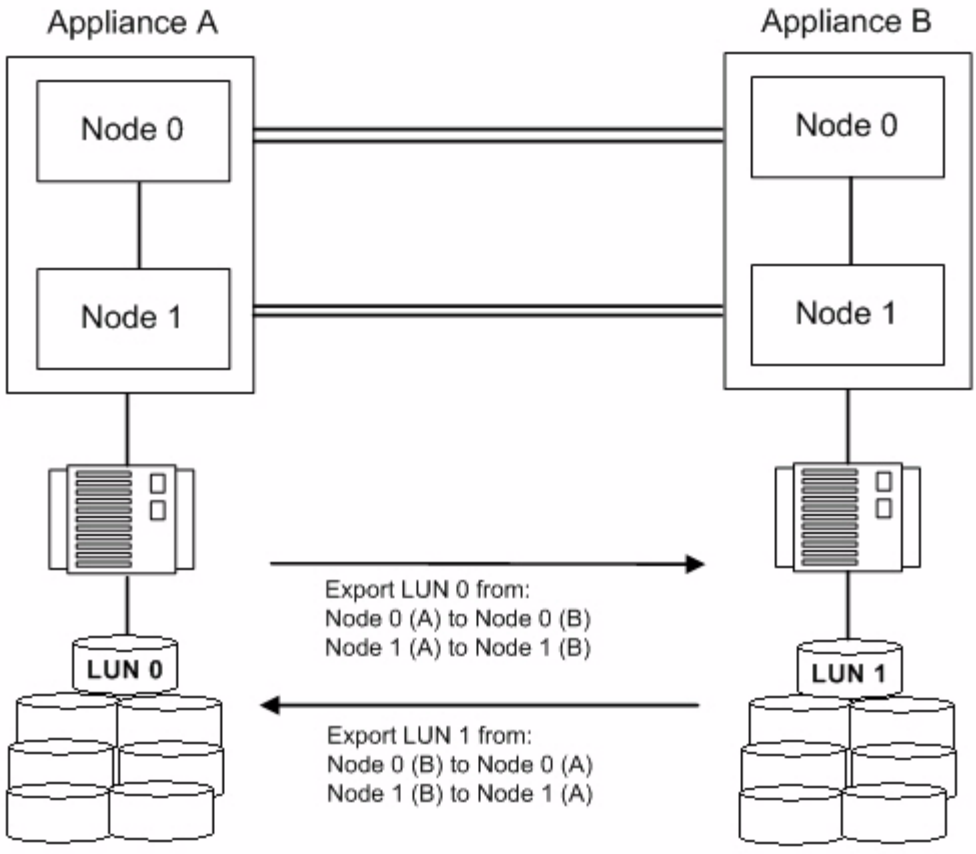

Figure 81: Exporting LUNs

Creating Exported LUNs

Note: This guide uses the terms "local appliance" and "remote appliance" to distinguish between the cascaded appliances. The local appliance is the appliance you are currently on.

To export a LUN, you must specify:

- **Host**—The remote node to which the LUN is being exported
- **LUN**—The LUN that you want to export to the remote node
- **Host LUN number**—The virtual number that identifies the exported LUN
- **Access control**—The type of access the remote node will have to the exported LUN: read only or read write
- **Target**—The target port of the remote node through which the LUN is exported

You must export the same LUN from both nodes of the appliance:

- Node 0 of Appliance A to Node 0 of Appliance B
- Node 1 of Appliance A to Node 1 of Appliance B

Procedure

To export a LUN:

1. Choose **LUN Map > create lunmap**.

   The **create lunmap** window opens (Figure 82).
2. Choose the remote appliance node in the **select host** field.
3. Choose the LUN to be exported in the **select lun** field.
4. Enter a number to identify the exported LUN in the **select host lun number** field. The available range for this field is 1–255 (per target port).
5. Choose the target port in the **select target** field.
6. Choose **read only** or **read-write** in the **access control** field.
7. Click **submit**.

   The **view lunmaps** window opens (Figure 84).
8. Repeat steps 2 through 7 to export the same LUN from the appliance's peer node.

Figure 82 shows how to export LUN 0 being exported from Node 0 on Appliance B to Node 0 on Appliance A.

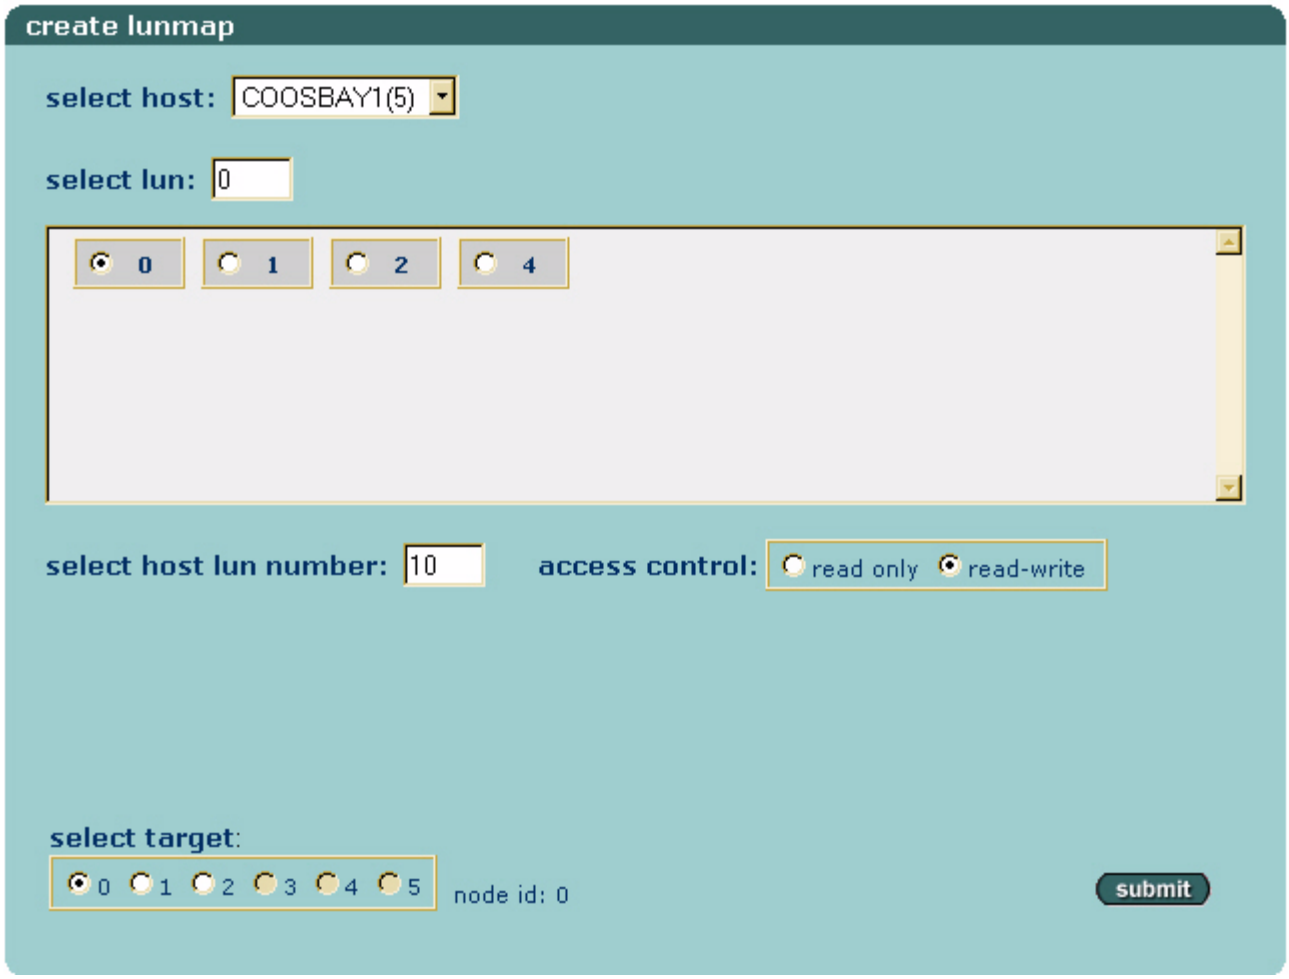

Figure 82: Exporting a LUN to Node 0 on Appliance A

Figure 83 shows how to export LUN 0 from Node 1 on Appliance B to Node 1 on Appliance A.

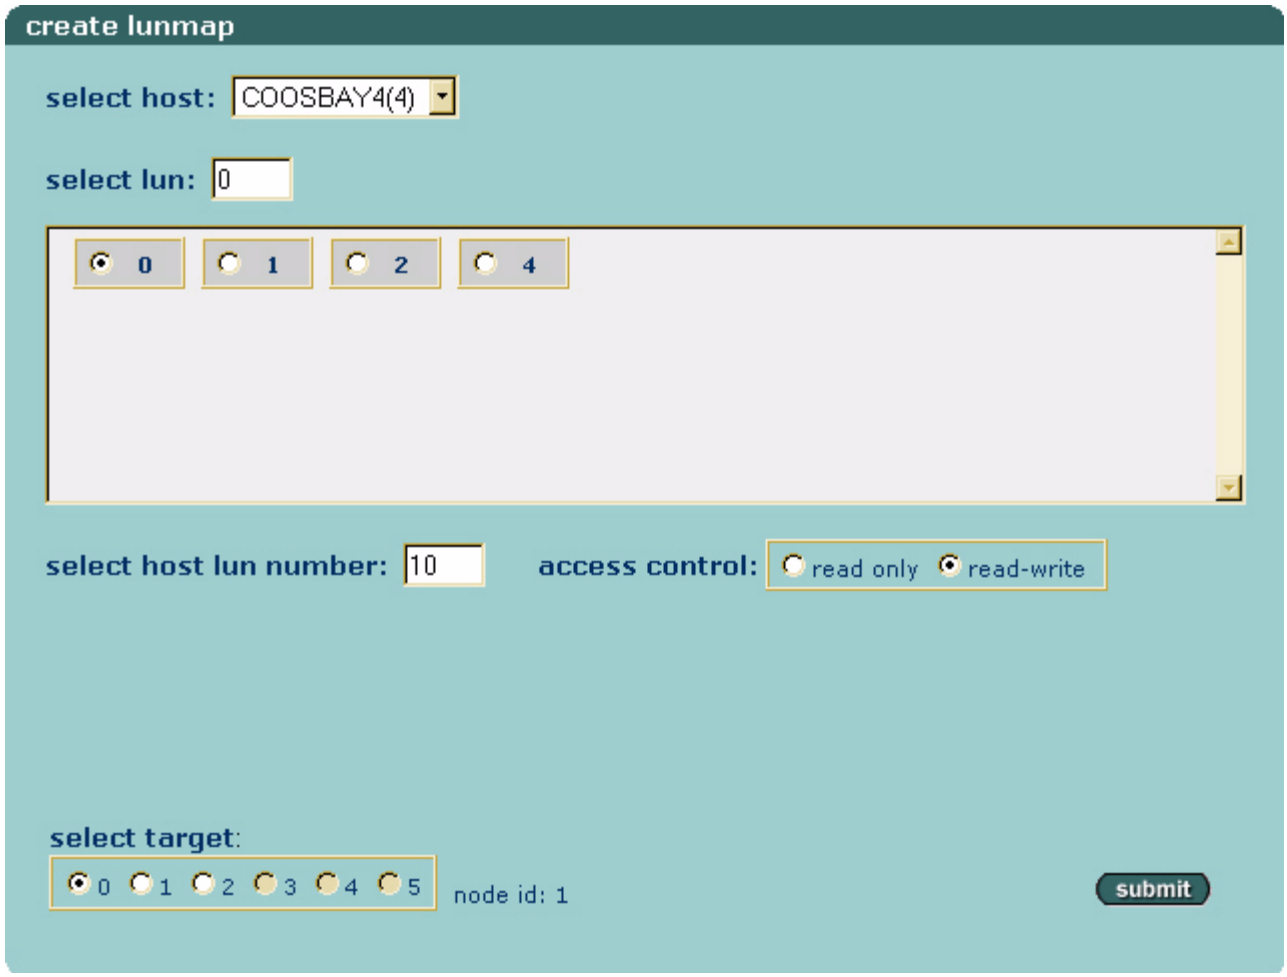

Figure 83: Exporting a LUN to Node 1 on Appliance A

Figure 84 shows the LUNs exported from Appliance A to Appliance B.

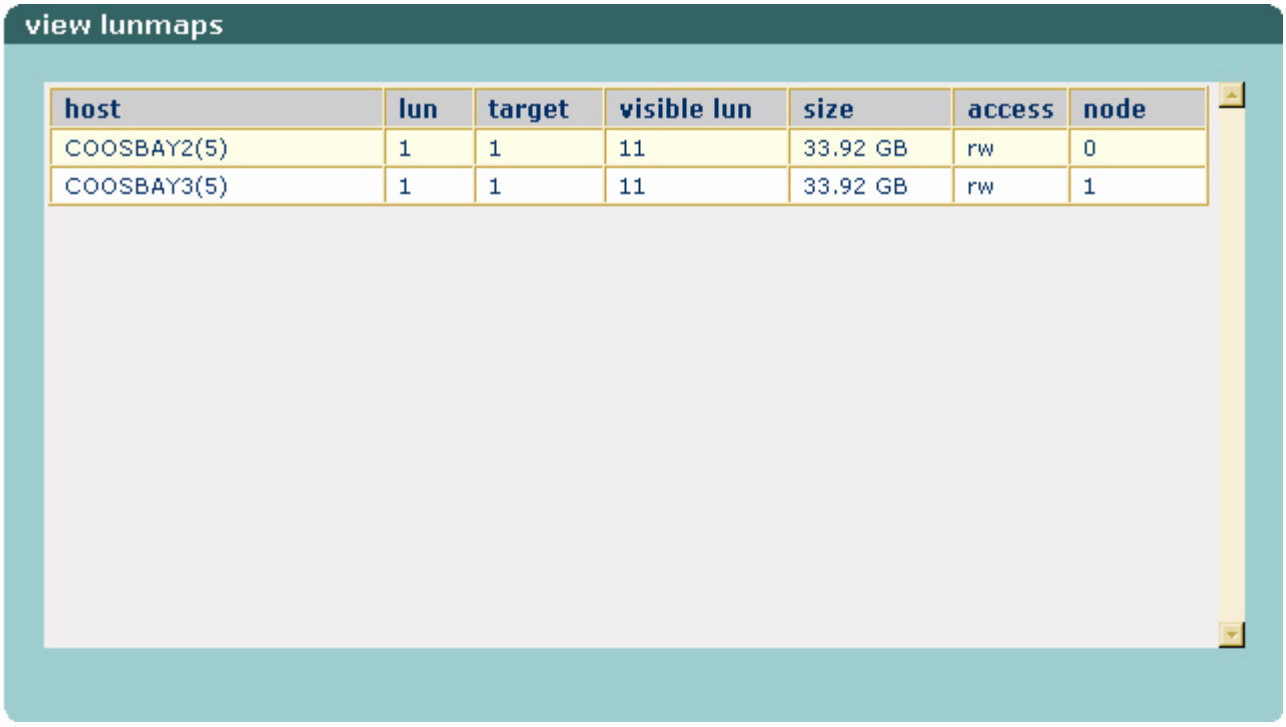

Figure 84: LUNs exported from Appliance A to Appliance B

Figure 85 shows the LUNs exported from Appliance B to Appliance A.

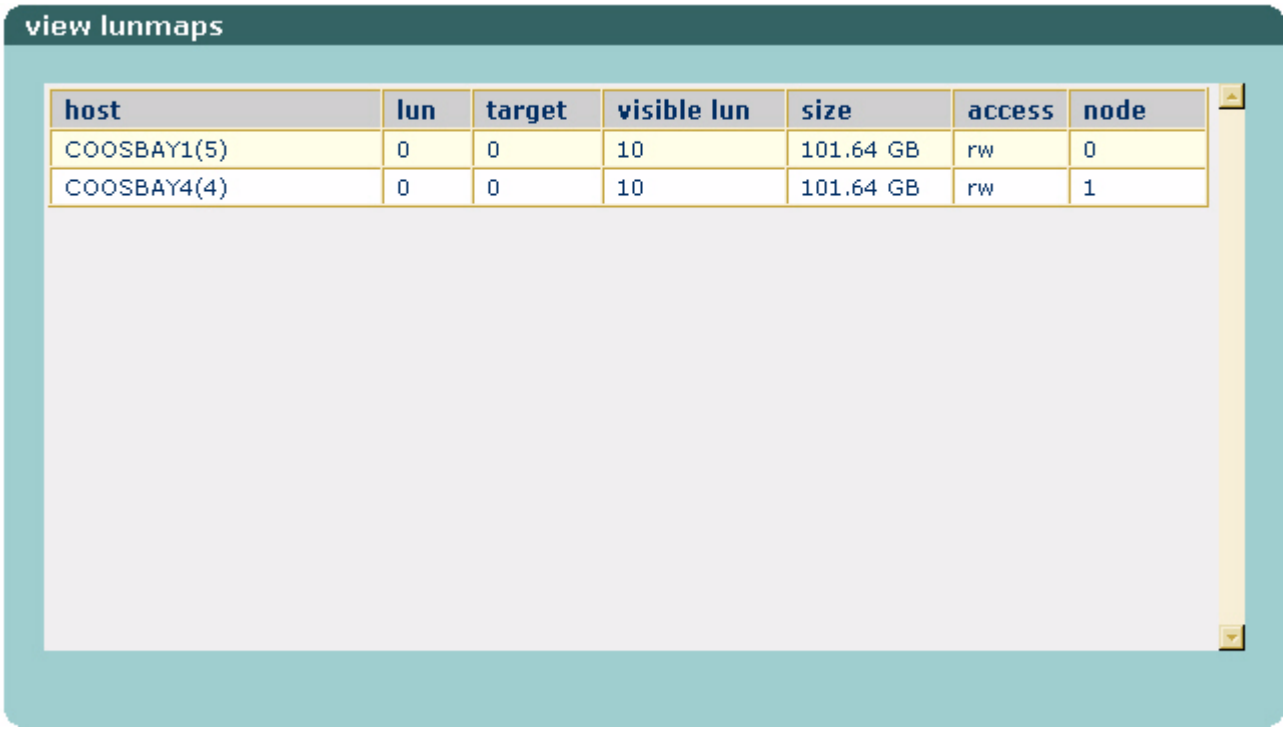

Figure 85: LUNs exported from Appliance B to Appliance A

Viewing Exported LUNs

You can view exported LUNs from the user interface and the command line interface.

User Interface

You can view exported LUNs using the following:

- LUN Matrix Window
- LUN Details Window

LUN Matrix Window

1. Select **LUNs > lun matrix**.

   The **view luns** window opens, displaying the matrix (Figure 86).

2. Use the scroll bar at the bottom of the window to view the last two columns, which are **exported lun** and **imported lun**. Columns marked with an "**x**" indicate how the LUN is being used:

   - **exported lun**—Any LUNs exported from the local appliance to the other cascaded appliance. For example, when you view the matrix from Appliance B, an "**x**" in the **exported lun** column indicates which LUNs you exported *to* Appliance A (Figure 86).
   - **imported lun**—Any LUNs exported from the other cascaded appliance to the local appliance. For example, when you view the matrix from Appliance B, an "**x**" in the **imported lun** column indicates which LUNs you exported *from* Appliance A (Figure 86).

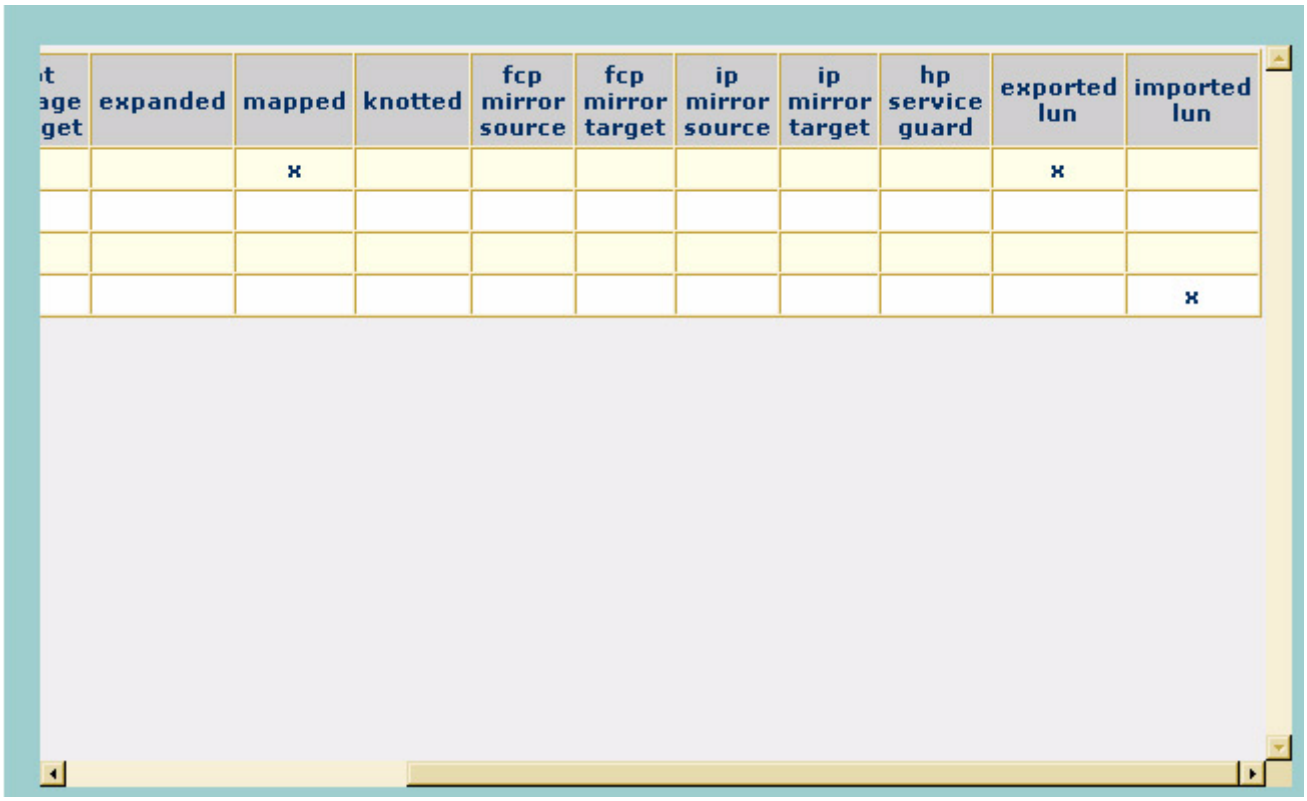

Figure 86: Viewing the LUN matrix from Appliance B

LUN Details Window

1. Choose **LUNs** from the menu bar.

   The **view luns** window opens.

2. Choose the LUN you want to view.
3. Click **details**.

   The **lun x** window opens.

4. Scroll to the bottom of the window for exported LUN information. The information that is displayed depends on the appliance and the node from which you are viewing the LUN. For example, when you view LUN 1 from Node 0 on Appliance A, the **lun x** window displays the following information (Figure 87):

   - **imported**—Displays **true** for a LUN that has ben exported to another appliance.
   - **exported**—Displays **true** for a LUN that has been exported from another appliance.
   - **imported as**—Indicates how the LUN is identified by the appliance to which it has been exported.

     This number is based on the total number of LUNs connected to the appliance. For example, Appliance B has three LUNs connected to it. When LUN 1 is exported to Appliance B from Appliance A, the LUN is identified as 4 because that is the next available number on Appliance B.

     Note: If the **imported** field displays **true**, this field name is **imported as**. When the **exported** field displays **true**, this field name is **exported as**.

   - **imported by ip O**—This field is not currently used.
   - **imported by host**—Identifies the node by its host number, which the appliance software generates when the nodes are cascaded to another appliance. The software generates the host number in the order in which each host is detected. For example, Node 0 of Appliance B is identified as Host 3 by Node 0 on Appliance A (Figure 87).

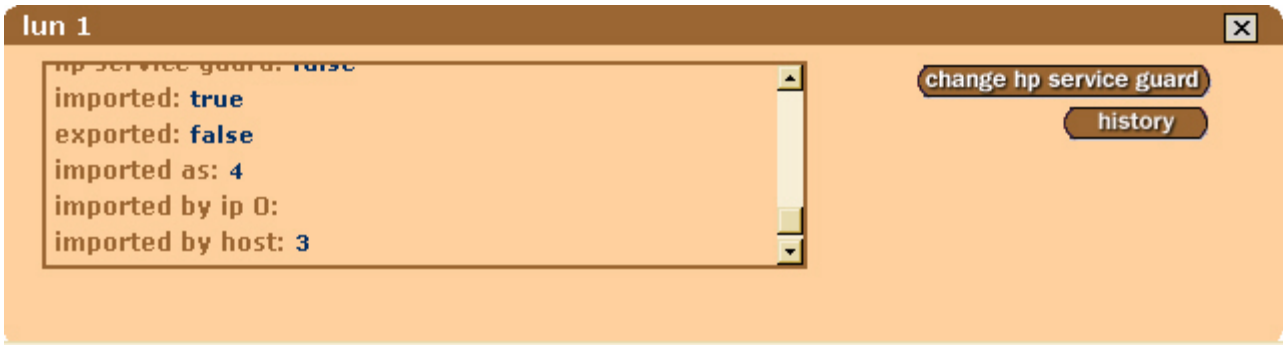

Figure 87: Viewing exported LUN details from Node 0 on Appliance A

When you view the details of LUN 1 from Node 1 on Appliance A, the only difference is the host number in the **imported by host** field (Figure 88). Node 1 on Appliance A identifies Node 1 on Appliance B as Host 4.

Figure 88: Viewing exported LUN details from Node 1 on Appliance A

When you view LUN 1 on Appliance B (the LUN exported from Appliance A), you view it as LUN 4, the number by which Appliance B identifies the LUN. The LUN details are displayed differently (Figure 89).

The **exported** field displays **true** because the LUN was exported from Appliance A. The **exported as** field displays **1** because the LUN is identified as LUN 1 on Appliance A. The **exported by host** field displays 1, indicating that Node 0 on Appliance B identifies Node 0 on Appliance A as Host 1.

Figure 89: Viewing imported LUN details from Node 0 on Appliance B

Command Line Interface

To view exported LUNs from the CLI, enter one of the following commands from the local appliance:

`get lun -exported`—View the LUNs you **exported to** another appliance.

`get lun -imported`—View the LUNs you **exported from** another appliance.

For example, enter the `get lun -exported` command from either node on Appliance A. The system displays the following output:

```
LUN        Size
1          33.92 GB
```

Enter the `get lun -imported` command from either node on Appliance A. The system displays the following output:

```
LUN       Size
19        101.64 GB
```

Cross-Appliance Mirror Details

To create a cross-appliance mirror, you must specify the following:

- **Source LUN**—The storage disk containing the data you want to replicate. The LUN you designate as the source of the mirror must be the LUN you exported *to* the remote appliance. For example, LUN 0 is exported from Appliance A to Appliance B, and LUN 1 is exported from Appliance B to Appliance A. If you want Appliance A to be the source of the mirror, LUN 0 must be the source LUN. If you want Appliance B to be the source of the mirror, LUN 1 must be the source LUN.
- **Target LUN**—The storage disk on which data will be replicated. The LUN you designate as the target of the mirror must be the LUN you exported *from* the remote appliance. For example, LUN 0 is exported from Appliance A to Appliance B. If you create the mirror from Appliance A, LUN 0 cannot be the target LUN. The target LUN must be one of the LUNs exported from Appliance B to Appliance A, such as LUN 1.
- **State**—The state of the source LUN and target LUN when you create the mirror. Choose **full resync** to ensure that the data on both LUNs is the same before replication begins. Choose **no resync** to enable replication immediately. Note, however, that existing data on the source LUN and target LUN may not be the same.
- **Allow target failover on remote site**—This option enables target failover on the remote appliance if you map the source LUN to hosts from the remote target appliance. Hosts are not notified of the source appliance failure and continue to access and write data to the source LUN. After you fix or replace the source appliance, unmap the LUN (from the host to the remote appliance) and perform a reverse resynchronization.

  Do not use this option if you want to use the local target LUN for an IP mirror. See "Using Exported LUNs for IP Mirrors" on page 136 for more information.

The Continuous Access Storage Appliance performs cross-appliance FCP mirroring in synchronous mode. Before a replication request is flagged as successful, a write success message must be returned from both the source and target LUNs.

Using Exported LUNs for IP Mirrors

The LUNs that you create for cross-appliance mirrors can also be used for IP mirrors. The procedure for creating the IP mirror remains the same but the following rules apply:

- You must create the cross-appliance FCP mirror before you create an IP mirror.
- You can only use the local source or the local target as the source LUN for an IP mirror.

See the chapter "Managing IP Mirrors" on page 147 for detailed information on creating IP mirrors.

Creating Cross-Appliance FCP Mirrors

To create a cross-appliance FCP mirror:

1. From the node of the local appliance, choose **mirrors > cams > create cam**.

   The **create cross appliance mirror** window opens (Figure 90).
2. Choose the source LUN in the **select source lun** field.
3. Choose the target LUN in the **select target lun** field.

   If there are no target LUNs to choose on this window, ensure that you have exported the LUNs correctly.
4. Choose **no resync** or **full resync** in the **select state** field.
5. Choose the **allow target failover on remote site** option if you have a non-clustered environment.
6. Click **submit**.

   The **view cross appliance mirrors** window opens, displaying the mirror you created (Figure 91).

Figure 90 shows the mirror you created from Appliance A. The source LUN is 1 and the target LUN is 19. If you create the mirror from Appliance B, the source LUN is 0 and the target LUN is 4.

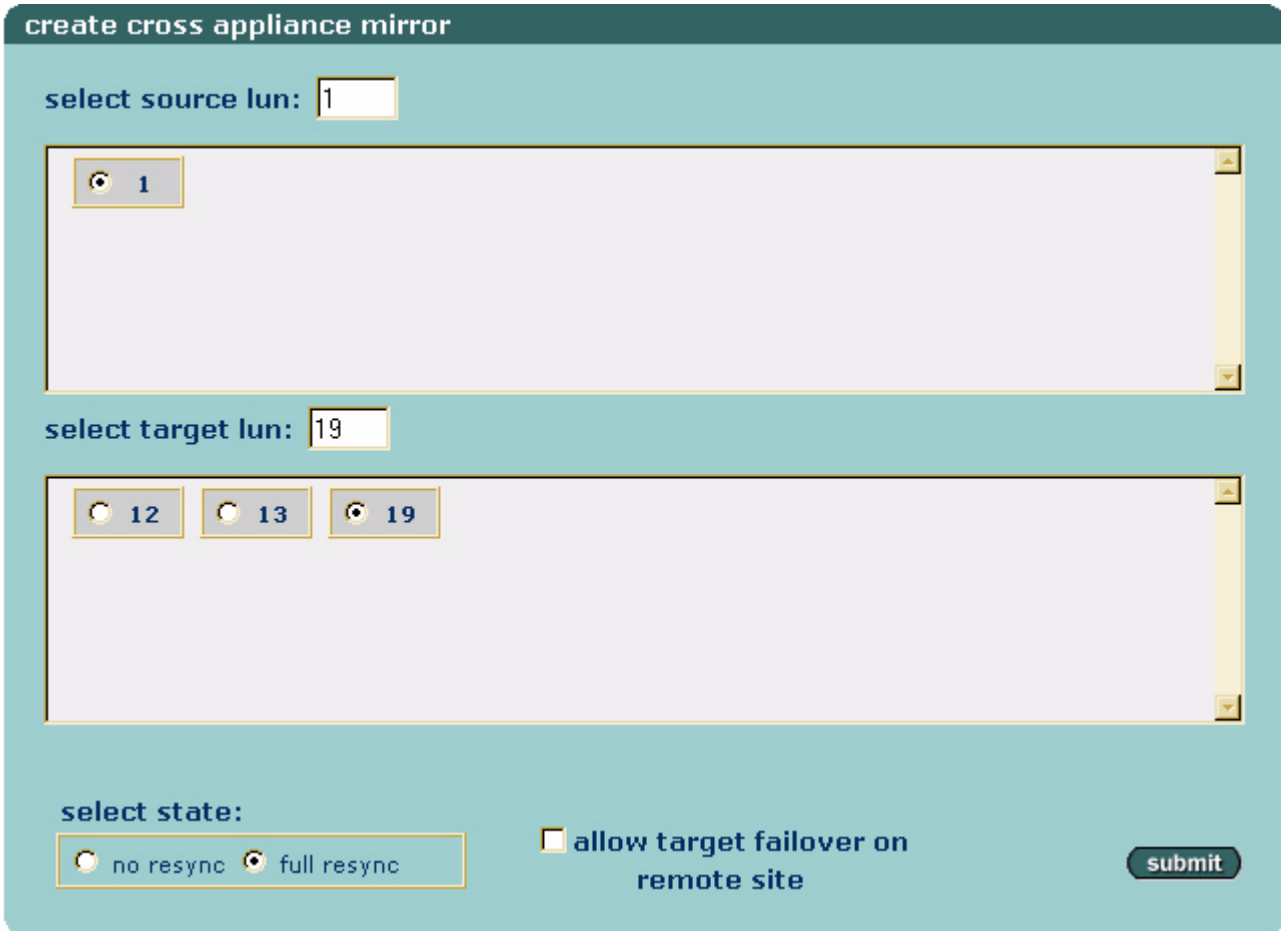

Figure 90: Create cross appliance mirror window

Table 7 describes the fields in the **view cross appliance mirrors** window (Figure 91 and Figure 92).

Table 7: Fields in the view cross appliance mirrors window

Field	Description
local source	The source LUN of the mirror as it is identified on the local appliance from which it was exported.
remote source	The source LUN as it is identified on the appliance to which it has been exported.
local target	The target LUN of the mirror as it is identified on the local appliance from which the LUN was exported.
remote target	The target LUN of the mirror as it is identified on the appliance to which the LUN has been exported.
state	The state of the data on the remote target LUN. Values are **Valid**, **Invalid**, or **Out of date**. A LUN goes into the **Out of date** state when the LUN is down. In this case, data on the LUN is still block consistent, but it is not current. A LUN goes into the **Invalid** state when a resynchronization has started, which causes the LUN to be block inconsistent. This means that the blocks of data on the source and target LUNs will be different during the time that the data is being copied from one LUN to the other. If a site disaster or temporary outage occurs on the source site appliance and the target LUN was down at the time, the target LUN on the target appliance may be in an **Out of date** or **Invalid** state.
block target	The state of the block on the remote target LUN.
imported lun ip	The IP address of the remote node from which the LUN was exported.
exported lun ip	The IP address of the remote node from which the LUN was exported.

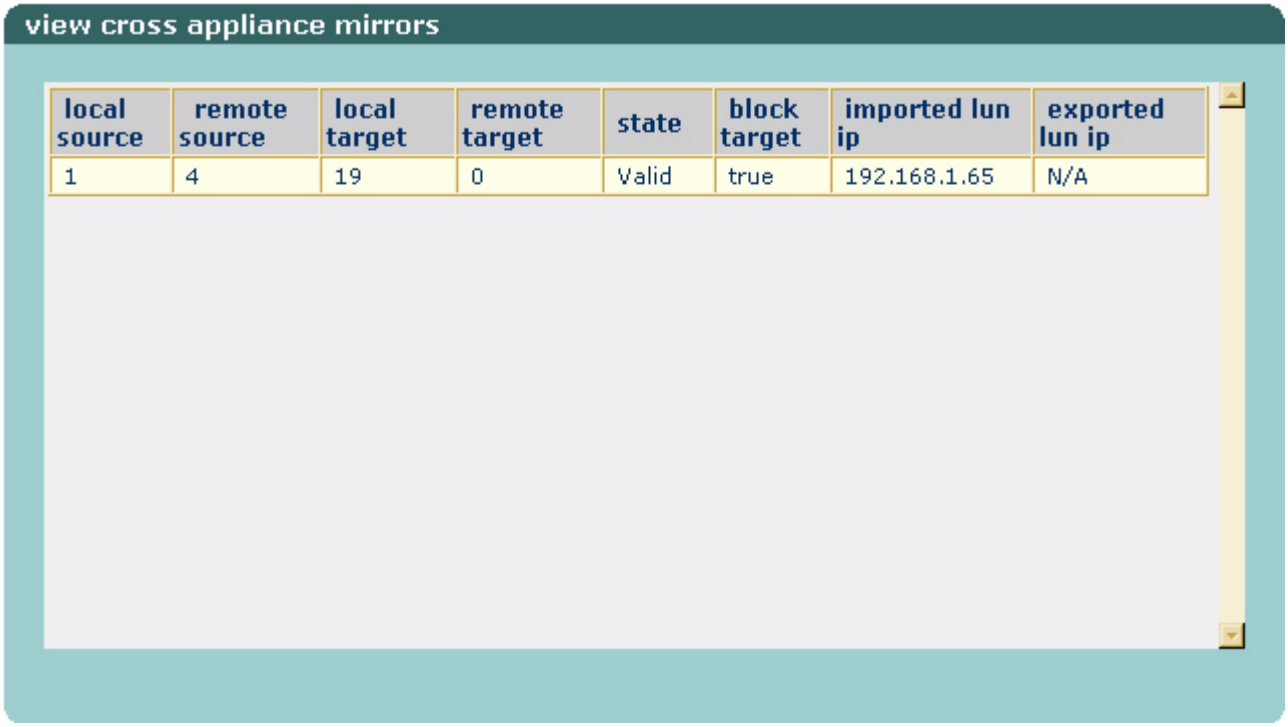

Figure 91: View cross-appliance mirrors window from Node 0 on Appliance A

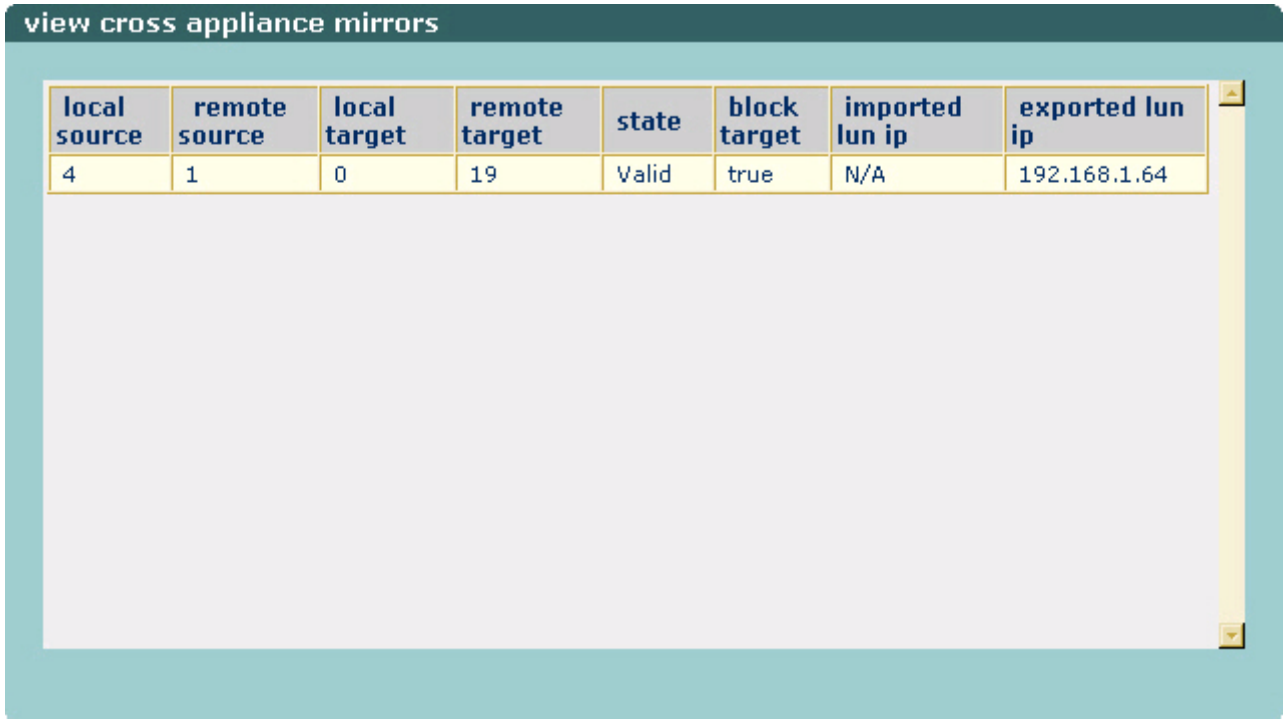

Figure 92: View cross-appliance mirrors window from Node 0 on Appliance B

About Mirror States

The mirror states for local FCP mirrors apply to cross-appliance FCP mirrors. See “About Mirror States” on page 114 for more information.

Pausing Cross-Appliance FCP Mirrors

Pausing a cross-appliance FCP mirror temporarily stops data replication between the source LUN and the target LUN. When you pause the mirror, the appliance software tracks new replication requests in the mirror's journal file.

Note: Disk redundancy is disabled for paused mirrors.

You can pause a mirror to save an image of the data on the target LUN. After you back up the data, resume mirroring.

To pause a cross-appliance FCP mirror:

1. From the local appliance, choose **mirrors > cams > pause-resync**.

   The **pause resync cross appliance mirrors** window opens (Figure 93).

2. Choose the mirror you want to pause.

   The **selected mirror** window opens (Figure 94).

3. Choose **pause** and click **submit**.

   The **pause resync cross appliance mirrors** window opens, showing the mirror in the Paused state (Figure 95).

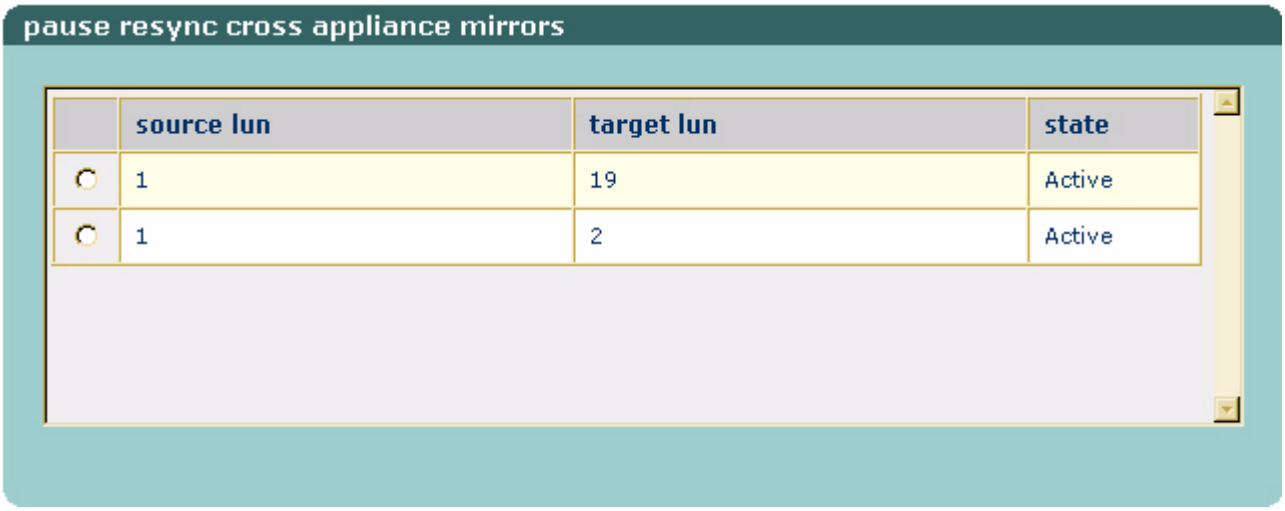

Figure 93: Pause resync cross appliance mirrors window

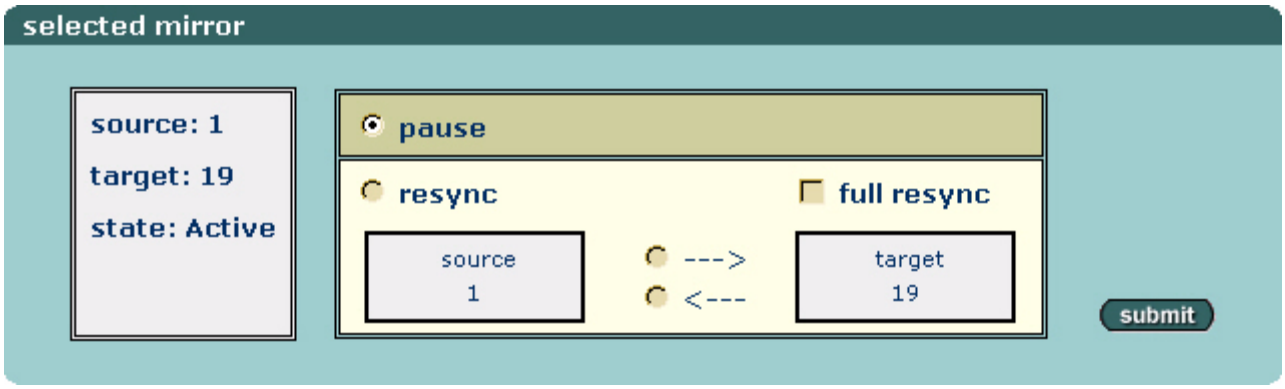

Figure 94: Selected mirror window

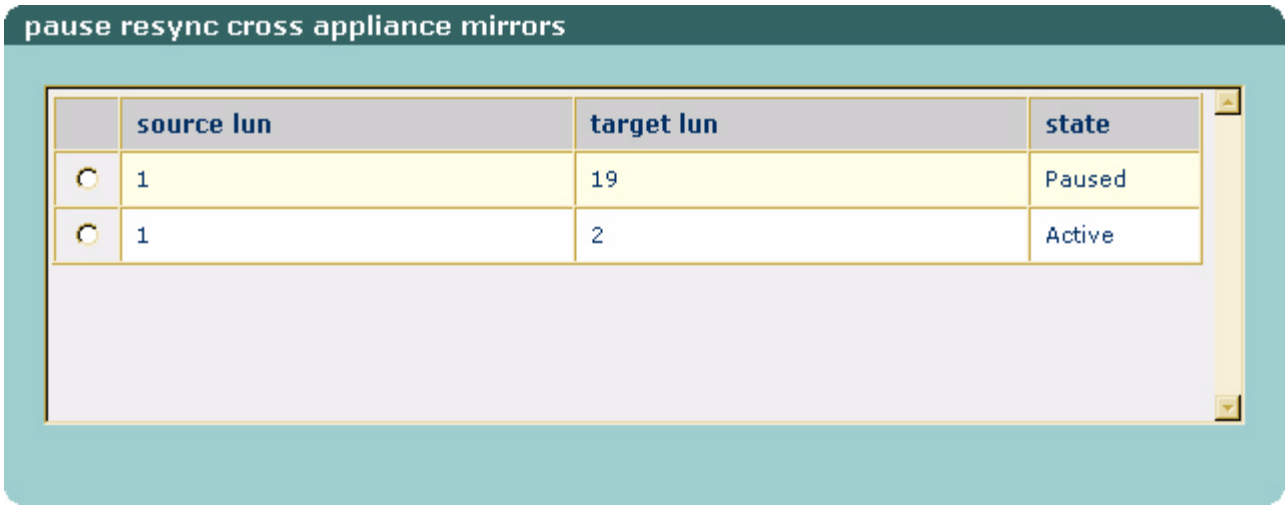

Figure 95: Pause resync cross appliance mirrors window (Paused mirror state)

Resynchronizing Cross-Appliance FCP Mirrors

Mirroring can stop if:

- A failure occurs on the local or remote appliance. (See "About Mirror States" on page 140.)
- You pause the mirror. (See "Pausing Cross-Appliance FCP Mirrors" on page 141).

Before you resume mirroring, you must resynchronize the data on the source and target LUNs so that both LUNs contain the same data. You must choose:

- **Data to be resynchronized**—Choose **resync** to use the data received since you paused the mirror. Choose **full resync** to use all data since you created the mirror.
- **Direction of the resynchronization**—Choose forward resynchronization (source to target) to update the target LUN with the source LUN's data. Choose reverse resynchronization (target to source) to update the source LUN with the target LUN's data. In either direction, the source LUN does the work—it sends the data (forward resynchronization) or obtains the data (reverse resynchronization).

To resynchronize a cross-appliance FCP mirror:

1. From the local appliance, choose **mirrors > cams > pause-resync**.

   The **pause resync cross appliance mirrors** window opens (Figure 95).

2. Choose the mirror you want to resynchronize.

   The **selected mirror** window opens (Figure 96).

3. Choose **resync** or **full resync**.
4. Choose the resynchronization direction–from source to target (forward) or from target to source (reverse).
5. Click **submit**.

   The **pause resync cross appliance mirrors** window opens, showing the mirror in the Resynchronizing state (Figure 97).

Note: When you view this mirror from the peer node, the mirror state is Waiting.

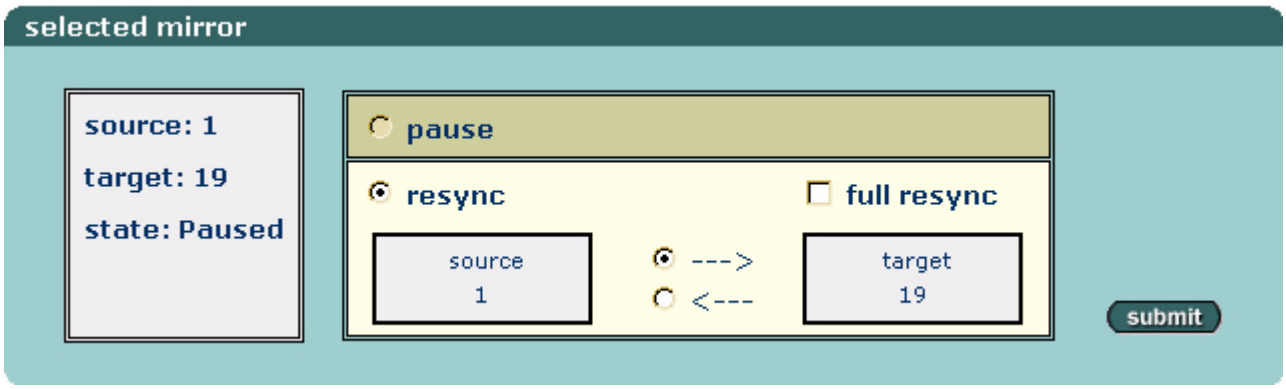

Figure 96: Selected mirror window

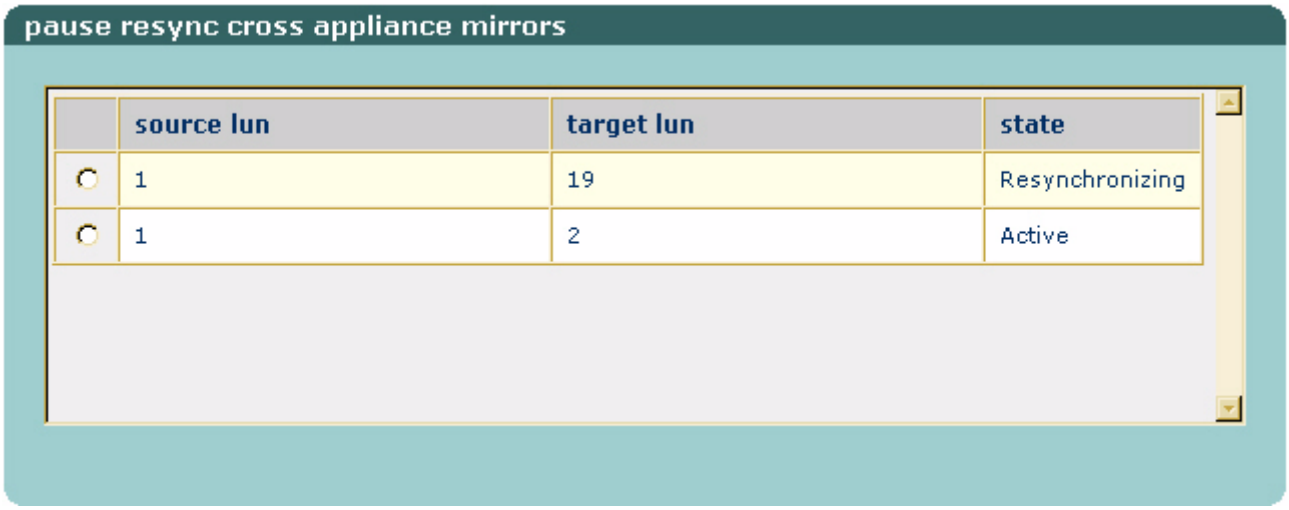

Figure 97: Pause resync cross appliance mirrors window (Resynchronizing mirror state)

Deleting Cross-Appliance FCP Mirrors

You delete a cross-appliance FCP mirror when you no longer need to replicate data.

After you delete the mirror on the local appliance, go to the remote appliance and unmap the LUNs exported to the local appliance. See "Deleting LUN Maps" on page 67 for more information.

To delete a cross-appliance FCP mirror:

1. From the local appliance, choose **mirrors > cams > delete cam**.

   The **delete cross appliance mirror** window opens (Figure 98).
2. Choose the mirror you want to delete.
3. Click **submit**.

   The **view cross appliance mirrors** window opens. The mirror you deleted is not listed.
4. Repeat steps 1 through 3 on the remote appliance.

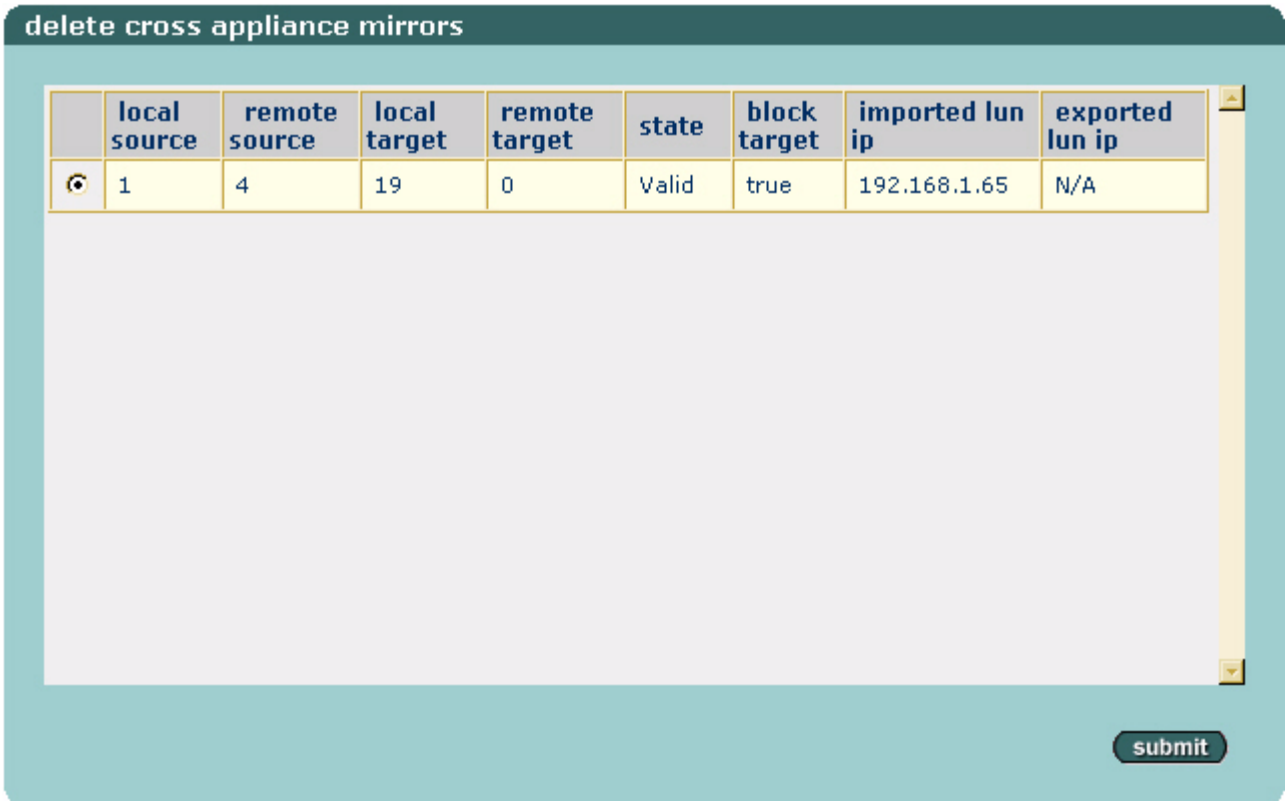

Figure 98: Delete cross appliance mirrors window

Managing IP Mirrors

11

This chapter defines IP mirroring and describes how to implement it in your network environment. It covers the following topics:

- About IP Mirrors, page 148
- Identifying Mirror Relationships, page 149
- Adding Mirror Nodes, page 156
- Adding Multiple Mirror Nodes, page 158
- IP Mirror Components, page 160
- Creating IP Mirrors, page 162
- About Mirror States, page 165
- Adjusting IP Mirror Parameters, page 173
- Pausing IP Mirrors, page 174
- Resynchronizing IP Mirrors, page 177
- Deleting IP Mirrors, page 182
- Removing Mirror Nodes, page 183

About IP Mirrors

IP mirroring is data replication between two appliances connected through a Gigabit Ethernet Internet Protocol (IP) link. The IP link enables data replication between appliances that are geographically dispersed. For example, you can mirror data from a storage device connected to an appliance in New York to another storage device connected to an appliance in London.

IP mirroring enables you to maintain a complete, identical image of data on disks in two locations. If an error disables one of the appliances or a site failure occurs, you can forward network traffic to the operational appliance and storage device without losing data access. After repairing or replacing the disabled appliance, copy data from the working image to the image connected to the newly operational appliance.

Up to three source appliances can replicate data to one target appliance, providing a common failover site for companies with multiple data centers. In this chapter, most examples show one source appliance replicating data to one target appliance.

Note: The target LUN of an IP mirror must be the same size or greater than the source LUN.

Identifying Mirror Relationships

This section covers the following topics:

- About Mirror Relationships
- About Secondary Paths
- Dual WAN Configuration
- Central WAN Configuration

About Mirror Relationships

You must use three different networks to effectively implement IP mirroring on the appliance. Each appliance node has three network interface card (NIC) ports, which operate as follows:

- **LAN Interface**—The network connection between your local LAN and each peer node, also referred to as the Management Link. The LAN Interface allows you to connect to the appliance management software through your web browser.
- **InterLink**—The Gigabit Ethernet connection between the peer nodes, also referred to as the heartbeat. Peer nodes always communicate through this link.
- **MirrorLink**—The Gigabit Ethernet connection between the local and remote nodes. The MirrorLink is used for IP mirroring only. The network for the MirrorLink must be different from the networks for the LAN Interface and InterLink.

You must make the following connections between the local and remote nodes:

- Primary mirror relationship from the local node to the remote node through the MirrorLink. For example, Node 1 on Appliance A has a mirror relationship to Node 3 on Appliance B, and vice versa.
- Secondary mirror relationship between the local node and the remote node. Establishing a secondary path ensures that the replication request completes successfully if the primary network path fails.

About Secondary Paths

There are two types of secondary paths:

- **Proxy**—through the InterLink
- **Direct**—through the MirrorLink

The type of path you choose depends on your network configuration. In a dual appliance setup, two common network configurations are:

- Dual wide area network (WAN)
- Central wide area network

Dual WAN Configuration

In a dual WAN configuration, a dedicated WAN connects each mirror relationship between the local and remote nodes (Figure 99).

Each node has three different network connections:

- **L**—LAN Interface
- **I**—InterLink
- **M**—MirrorLink

The mirror relationships are:

- Node 1 to Node 3
- Node 3 to Node 1
- Node 2 to Node 4
- Node 4 to Node 2

You must create the mirror relationship in both directions (Node 1 to Node 3 and Node 3 to Node 1) to establish communication between the nodes. If you only define a mirror relationship from Node 1 to Node 3, Node 3 cannot communicate with Node 1 because you did not create the mirror relationship from Node 3 to Node 1.

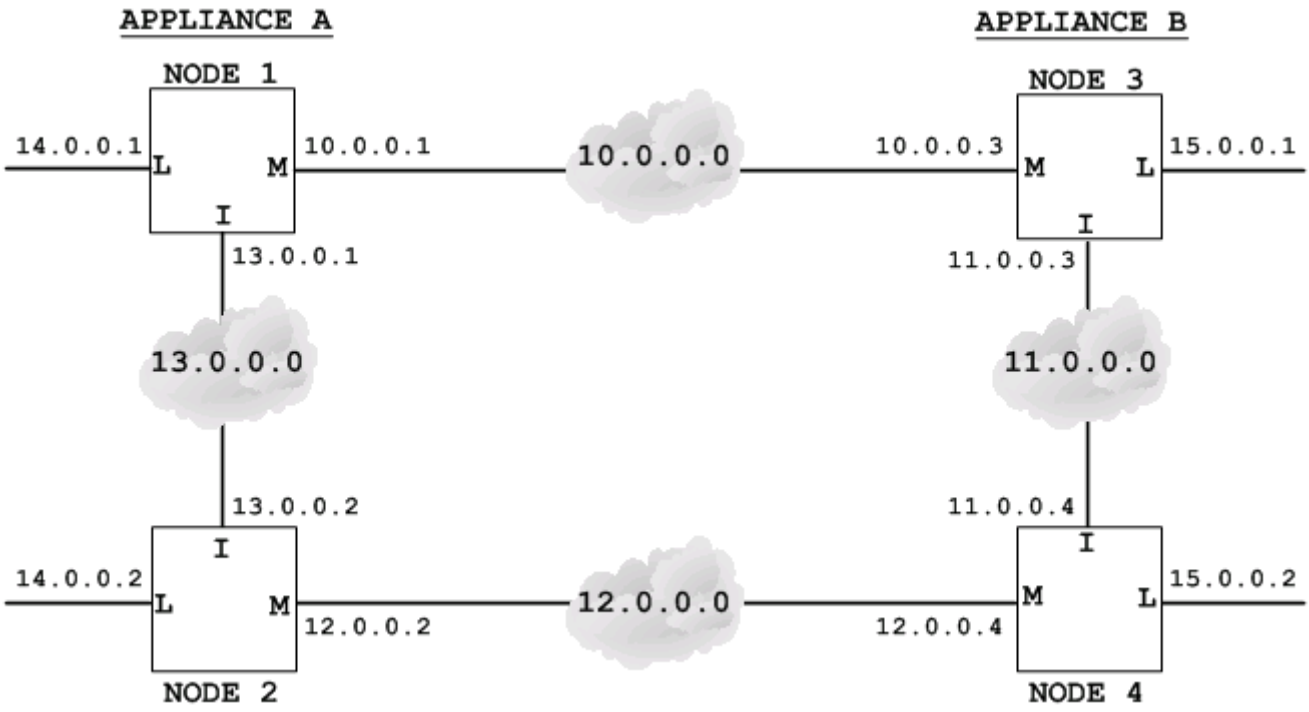

Figure 99: Dual WAN configuration

Using the Proxy Path

The primary mirror relationship is between Node 1 and Node 3. All data replication requests are sent through the primary path—the WAN connection between Node 1 and Node 3. If the nodes cannot communicate through the primary path, configure an alternate path to ensure that the request completes successfully. In a dual WAN configuration, use the proxy path as the alternate.

Note: Using the InterLink can impact performance. Use the proxy path *only* if the primary path fails.

The proxy path works as follows (Figure 100):

1. As a data replication request is processing, a failure occurs on the MirrorLink between Node 1 and Node 3.
2. The request traverses the InterLink between Node 1 and Node 2 on the local appliance.
3. The request crosses the MirrorLink between Node 2 and Node 4 to complete the request.

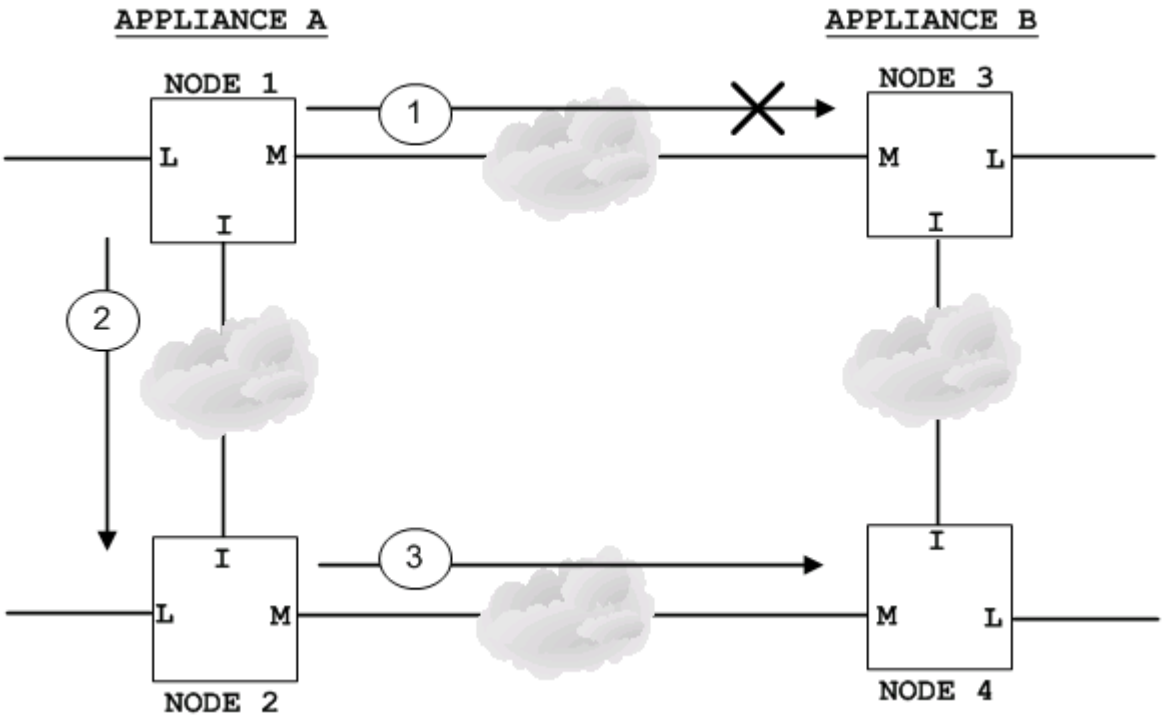

Figure 100: Proxy path example

Mirror Relationship Example

Establish the mirror relationships as follows:

- **Node 1**
  - Mirror node—10.0.0.3 (primary)
  - Peer node—13.0.0.2 (secondary/proxy)
- **Node 2**
  - Mirror node—12.0.0.4 (primary)
  - Peer node—13.0.0.1 (secondary/proxy)
- **Node 3**
  - Mirror node—10.0.0.1 (primary)
  - Peer node—11.0.0.4 (secondary/proxy)
- **Node 4**
  - Mirror node—12.0.0.2 (primary)
  - Peer node—11.0.0.3 (secondary/proxy)

Central WAN Configuration

In a central WAN configuration, one WAN connects all mirror relationships between the local and remote nodes (Figure 101).

Each node has three different network connections:

- **L**—LAN Interface
- **I**—InterLink
- **M**—MirrorLink

The mirror relationships are:

- Node 1 to Node 3
- Node 3 to Node 1
- Node 2 to Node 4
- Node 4 to Node 2

You must create the mirror relationship in both directions (Node 1 to Node 3 and Node 3 to Node 1) to establish communication between the nodes. If you only define a mirror relationship from Node 1 to Node 3, Node 3 cannot communicate with Node 1 because you did not create the mirror relationship from Node 3 to Node 1.

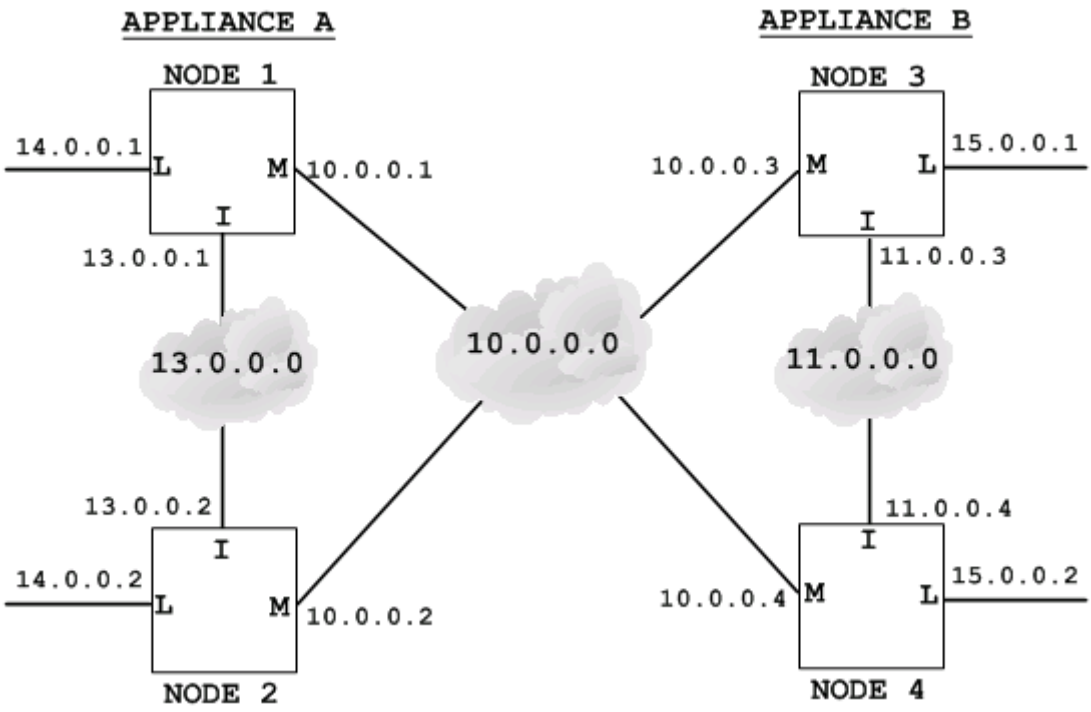

Figure 101: Central WAN configuration

Using the Direct Path

The primary mirror relationship is between Node 1 and Node 3. All data replication requests are sent through the primary path–the WAN connection between Node 1 and Node 3. If the nodes cannot communicate through the primary path, configure an alternate path to ensure that the request completes successfully. In a central WAN configuration, use either the proxy path or the direct path as the alternate. See “Using the Proxy Path” on page 151 for information about the proxy path.

The direct path works as follows (Figure 102):

1. As a data replication request is processing, a failure occurs on the MirrorLink between Node 3 and the WAN.
2. The request crosses the MirrorLink between Node 1 and Node 4 to complete the request.

The direct path provides the same performance as the primary path (both use the MirrorLink), but the direct path only works if the network problem is isolated between the WAN and Node 3. If the network problem occurs between Node 1 and the WAN, neither the primary path nor the direct path can work. If Node 1 cannot communicate with the WAN, it cannot communicate with either Node 3 (primary path) or Node 4 (direct path). Therefore, HP recommends that you use the proxy path in a central WAN configuration.

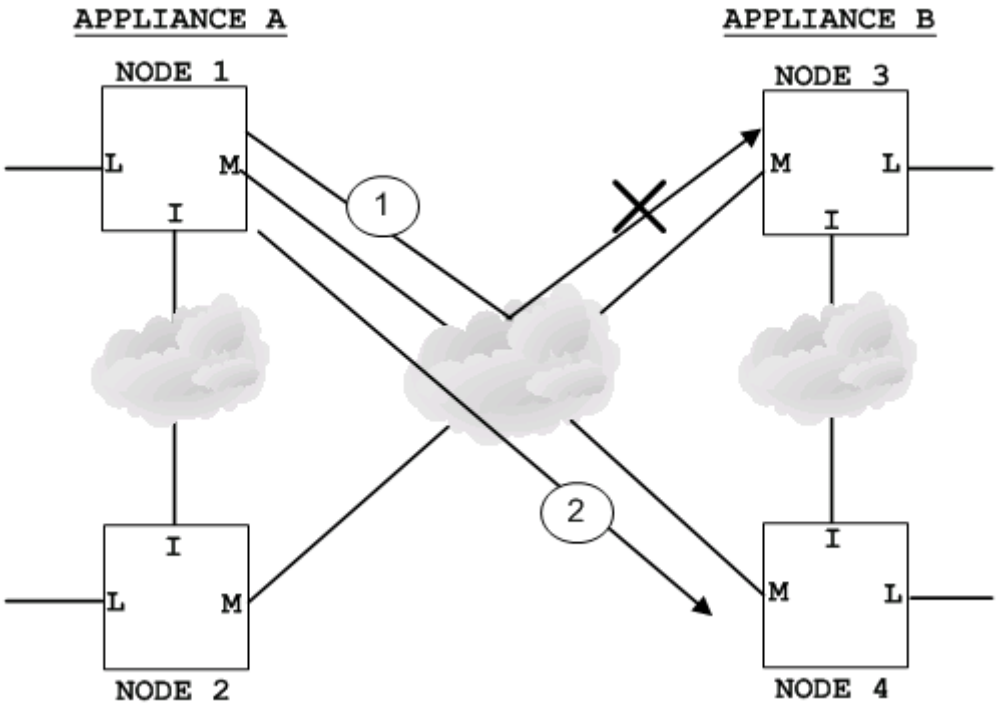

Figure 102: Direct path example

Mirror Relationship Example

Establish the mirror relationships as follows:

- **Node 1**
  - Mirror node—10.0.0.3 (primary)
  - Mirror node—10.0.0.4 (secondary/direct)

    OR
  - Peer node—13.0.0.2 (secondary/proxy)
- **Node 2**
  - Mirror node—10.0.0.4 (primary)
  - Mirror node—10.0.0.3 (secondary/direct)

    OR
  - Peer node—13.0.0.1 (secondary/proxy)
- **Node 3**
  - Mirror node—10.0.0.1 (primary)
  - Mirror node—10.0.0.2 (secondary/direct)

    OR
  - Peer node—11.0.0.4 (secondary/proxy)
- **Node 4**
  - Mirror node—10.0.0.2 (primary)
  - Mirror node—10.0.0.1 (secondary/direct)

    OR
  - Peer node—11.0.0.3 (secondary/proxy)

Adding Mirror Nodes

You establish mirror relationships by adding mirror nodes. Mirror nodes allow the local and remote appliances to communicate.

To add a mirror node, you must specify:

- **Mirror node**—The remote appliance node used for the primary mirror relationship.
- **Secondary node**—The appliance node used for the secondary mirror relationship. Use either the peer node of the local node (InterLink) or the other mirror node on the remote appliance (MirrorLink).
- **Alternate path option**—The type of path used for the secondary mirror relationship. Choose **proxy** when using the InterLink connection between the peer nodes (Node 1 and Node 2). Choose **direct** when using the MirrorLink connection between the local node (Node 1) and the peer node (Node 4) of the remote mirror node. See "Identifying Mirror Relationships" on page 149 for more information about alternate path options.

To add a mirror node:

1. Choose **nodes > add mirror node**.

   The **add mirror node** window opens (Figure 103).
2. Enter the MirrorLink IP address (or name) of the remote appliance node in the **mirror node** field.
3. Enter the IP address (or name) of the node to be used as the alternate path in the **secondary ip** field.
4. Do one of the following:
   - Choose **proxy** if you entered the InterLink IP address in the **secondary IP** field.
   - Choose **direct** if you entered the MirrorLink IP address in the the **secondary IP** field. The default option is **direct**.
5. Click **submit**.
6. Repeat steps 1 through 5 for all local and remote appliance nodes.

Note: If you entered the node name instead of the node IP address, ensure that you assigned a unique name to each node IP address during appliance installation.

Figure 103 shows how you create a mirror relationship from Node 3 to Node 1, using the central WAN configuration example from page 153.

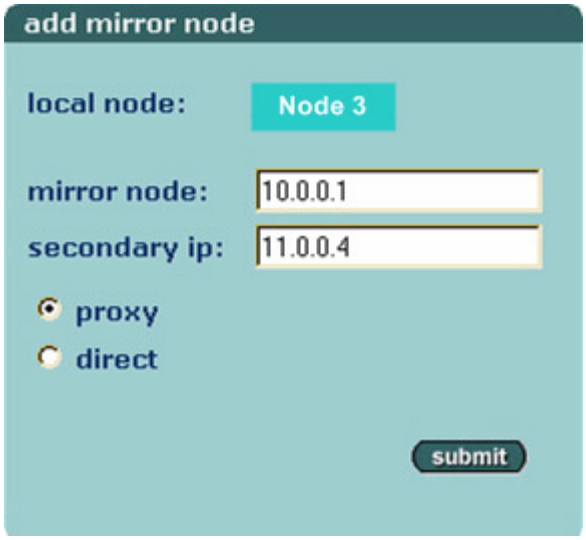

Figure 103: Add mirror node window

To create a mirror relationship from Node 3 to Node 1:

1. Enter the MirrorLink IP address for Node 1 in the **mirror node** field.
2. Enter the InterLink IP address for Node 4 in the **secondary ip** field.
3. Select **proxy** as the alternate path.
4. Click **submit**.

Adding Multiple Mirror Nodes

This section covers the following topics:

- About Many-to-One IP Mirroring
- Identifying Mirror Relationships
- IP Mirror Components

About Many-to-One IP Mirroring

The Continuous Access Storage Appliance supports a maximum of three source appliances replicating data to one target appliance. For example, Appliance A, Appliance B, and Appliance C can all replicate data to Appliance D (Figure 104).

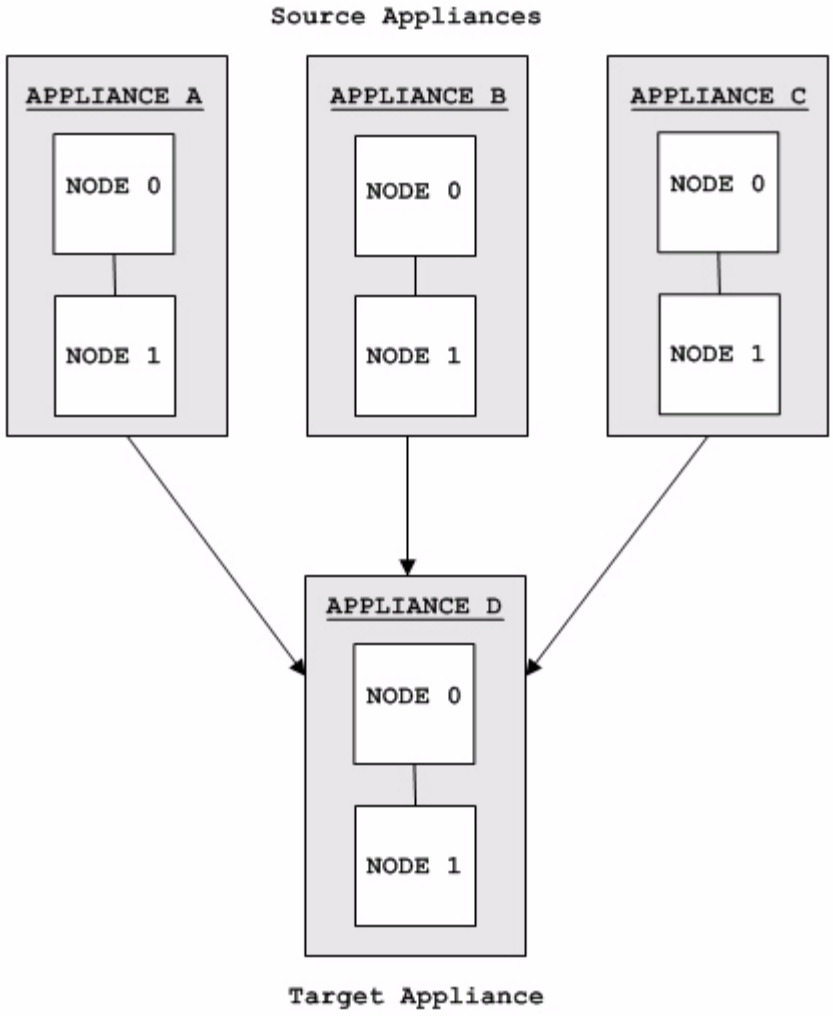

Figure 104: Three sources mirroring to one target

Identifying Mirror Relationships

This subsection covers the following topics:

- Sources to Target
- Target to Sources

Sources to Target

For each source (Appliances A, B, and C), the mirror relationships to Appliance D are as follows:

- From Node 0 of Appliance A, add Node 0 of Appliance D as a mirror node.
- From Node 1 of Appliance A, add Node 1 of Appliance D as a mirror node.
- From Node 0 of Appliance B, add Node 0 of Appliance D as a mirror node.
- From Node 1 of Appliance B, add Node 1 of Appliance D as a mirror node.
- From Node 0 of Appliance C, add Node 0 of Appliance D as a mirror node.
- From Node 1 of Appliance C, add Node 1 of Appliance D as a mirror node.

See "Adding Mirror Nodes" on page 156 for the steps to create these relationships.

Target to Sources

After you create the mirror relationship from each source to the target, you must create the relationship from the target to each source. The mirror relationships from Appliance D to each source are as follows:

- From Node 0 of Appliance D, add Node 0 of Appliance A as a mirror node.
- From Node 1 of Appliance D, add Node 1 of Appliance A as a mirror node.
- From Node 0 of Appliance D, add Node 0 of Appliance B as a mirror node.
- From Node 1 of Appliance D, add Node 1 of Appliance B as a mirror node.
- From Node 0 of Appliance D, add Node 0 of Appliance C as a mirror node.
- From Node 1 of Appliance D, add Node 1 of Appliance C as a mirror node.

Use the user interface to create the first set of mirror relationships to one source (for example, from Appliance D to Appliance A). For the remaining mirror relationships (Appliance D to Appliance B and Appliance D to Appliance C), you must use the `add node` command in the CLI. See "Adding Mirror Nodes" on page 156 for the procedure and the command syntax.

When you create the mirror relationships from the target to the first source, establish both a primary path and a secondary path. For the mirror relationships to the remaining sources, follow these guidelines:

- If you specified the proxy path for the first source, enter the primary path for the remaining sources only. Do not specify an alternate path.
- If you specified the direct path for the first source, you must specify a primary path and a secondary path for the remaining sources.

IP Mirror Components

This section covers the following topics:

- Synchronous Mirroring
- Asynchronous Mirroring

Mirror node relationships enable the local appliance and the remote appliance to communicate and to see the storage that is connected to each appliance.

To create an IP mirror, you must specify:

- **Source LUN**–The storage disk containing the data you want to replicate.
- **Target LUN**–The storage disk containing the replicated data.
- **Mode**–The manner in which replication requests are completed on the source LUN and target LUN. Options are synchronous (default) and asynchronous. See "Synchronous Mirroring" and "Asynchronous Mirroring" for more information.
- **State**–The state of the source LUN and target LUN at mirror creation. Choose **full resync** to ensure the data on both LUNs is the same before replication begins. Choose **no resync** to enable replication immediately; however, existing data on the source LUN and target LUN may not be the same.

Synchronous Mirroring

With synchronous mirroring, a data replication request must be completed on both the source LUN and the target LUN before the source LUN sends notification to the host. The request takes longer to complete, but it ensures that the data is exactly the same on both the source and target LUNs. HP recommends synchronous mirroring when you are replicating highly critical data.

Note: The appliance does not support bidirectional synchronous IP mirroring.

Asynchronous Mirroring

This subsection covers the following topics:

- About Guaranteed Write Ordering
- About the Drainer

With asynchronous mirroring, a data replication request must be completed on the source LUN before the source LUN sends notification to the host. The request completes more quickly, but the data on the source LUN and the target LUN is not the same until the request completes on the target LUN. HP recommends asynchronous mirroring when replicating data over great distances.

Note: The appliance supports bidirectional asynchronous IP mirroring.

About Guaranteed Write Ordering

All asynchronous replication requests are first sent to the asynchronous IP mirror (AIPM) queue on the local appliance in the order they are received (write ordering). The AIPM queue is a partition on the shared storage (connected to both appliance nodes) that is created automatically during appliance installation. The size of the queue is 23 GB (or 325K I/Os).

Replication requests may not reach the remote appliance because:

- Network communication to the remote appliance is disabled.
- The network cannot transfer requests as quickly as they are received.

About the Drainer

The drainer (an appliance software task) removes the requests from the AIPM queue and sends them to the remote appliance in the order they were received on the local appliance. By default, the drainer is activated.

You can stop the drainer so that all asynchronous mirroring tasks are completed during periods of low host I/O activity, which HP *strongly recommends*. Adding or deleting asynchronous mirrors during heavy host I/O activity may cause the appliance to respond more slowly to new management requests due to the extra processing the appliance is doing during heavy I/O. Additionally, you should always perform a resynchronization before enabling an asynchronous mirror. Resynchronization requires network resources that may also affect appliance responsiveness during heavy I/O.

You can activate and freeze the drainer from the user interface or the CLI. In the user interface, the **ip mirrors** submenu has two options:

- **activate drainer**
- **freeze drainer**

In the CLI, the `mod param` command includes the `-drainer` parameter. Refer to "Mirroring Commands" in the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide* for more information about the `mod param` command.

If the AIPM queue becomes full, any additional replication requests are tracked in the source LUN's journal files and you must resynchronize the mirror. The only time write ordering cannot be guaranteed is during an asynchronous mirror resynchronization.

When the queue is eventually emptied, the appliance begins to replicate data from the journal file and the mirror enters the *resynchronization* state. It remains in this state until the transfer of data associated with the journal is complete. See "About Mirror States" on page 165 for more information about asynchronous mirror states.

Moving the Queue

You can move the AIPM queue from the appliance's shared storage to another storage device that is connected to the appliance. Moving the AIPM queue to another storage device provides more space in the queue, if you require it.

You can only move the AIPM queue using the CLI. The command for moving the queue is `mod lun -lun LUN -setq`. Refer to the "MOD LUN" section on page 60 of the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide* for more information about this command.

Creating IP Mirrors

To create an IP mirror:

1. Choose **mirrors > ip mirror > create ip mirror**.

   The **create ip mirror** window opens (Figure 105).
2. Choose the source LUN in the **select source lun** field.
3. Choose the target LUN in the **select target lun** field.
4. Choose **sync** or **async** in the **select mode** field.
5. Choose **no resync** or **full resync** in the **select state** field.
6. Click **submit**.

   The **view ip mirrors** window opens. Figure 106 shows the mirror from the source appliance; Figure 107 shows the mirror from the target appliance.

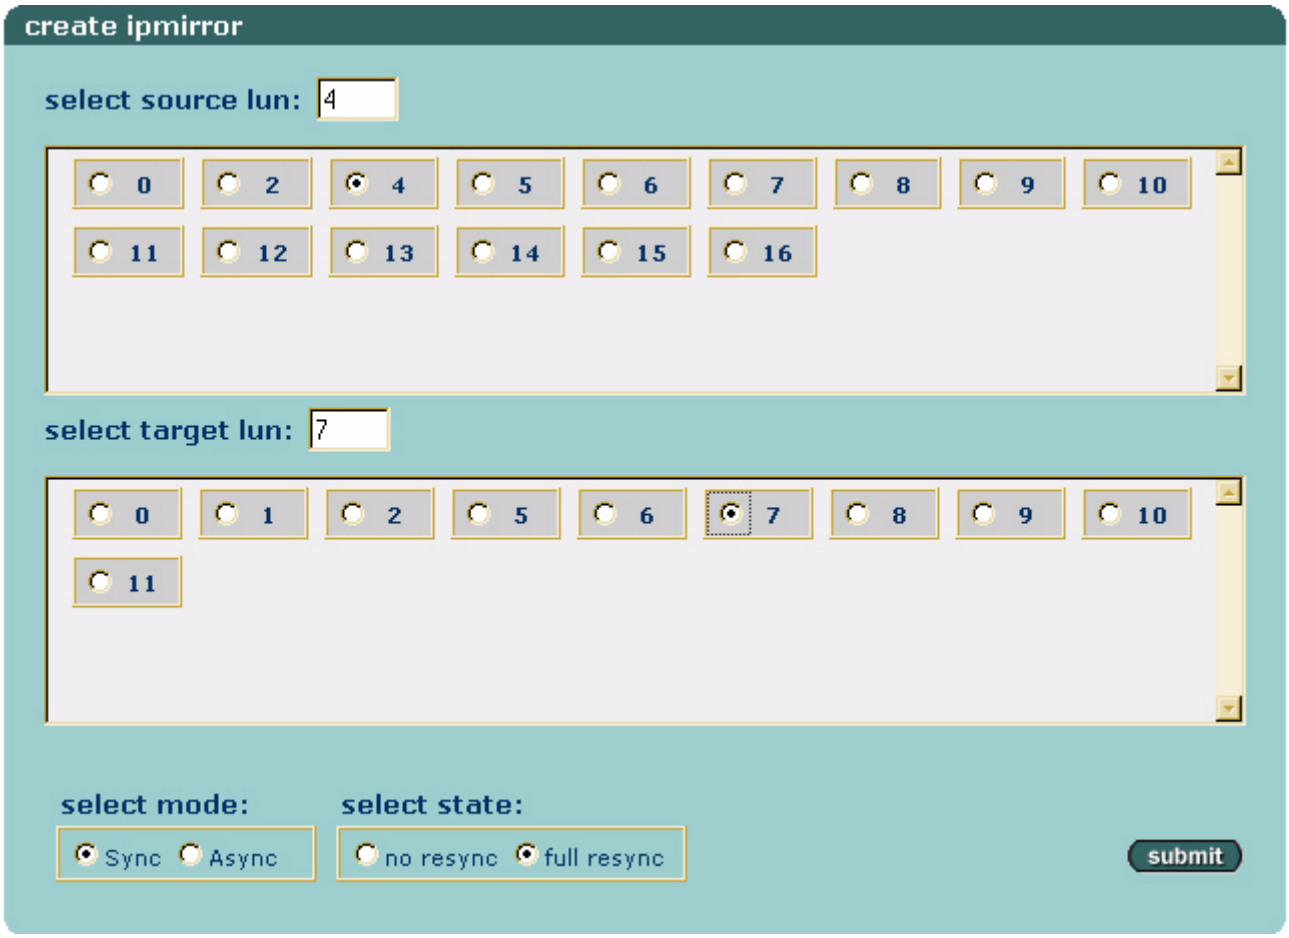

Figure 105: Create ip mirror window

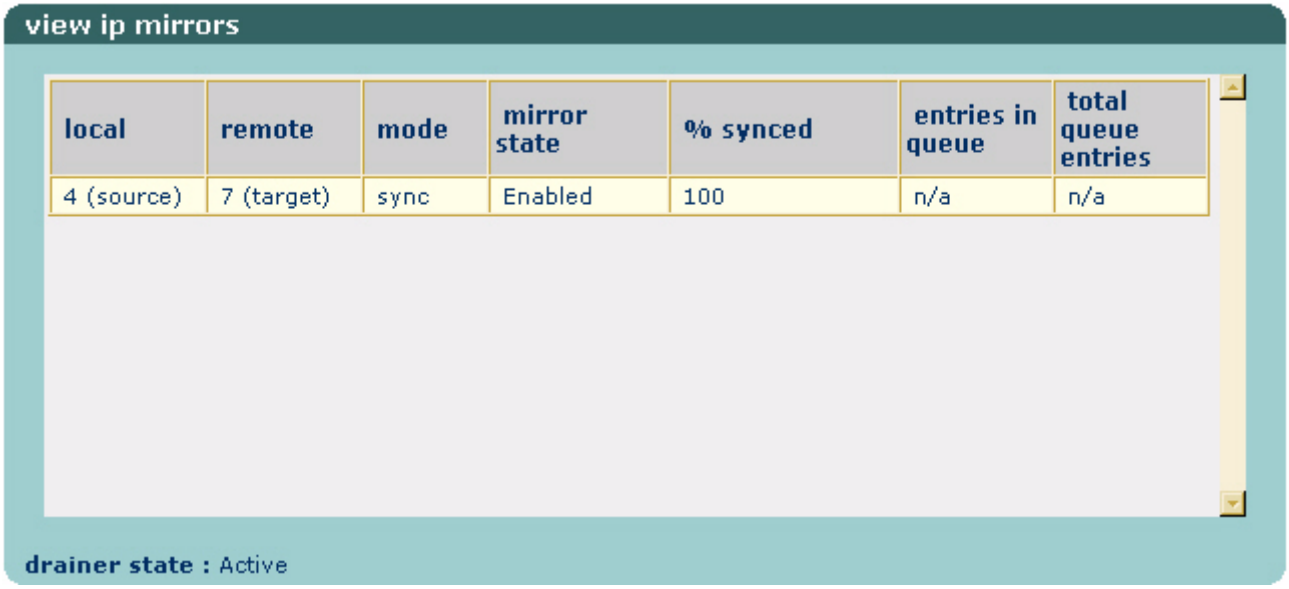

Figure 106: View ip mirrors window (from source appliance)

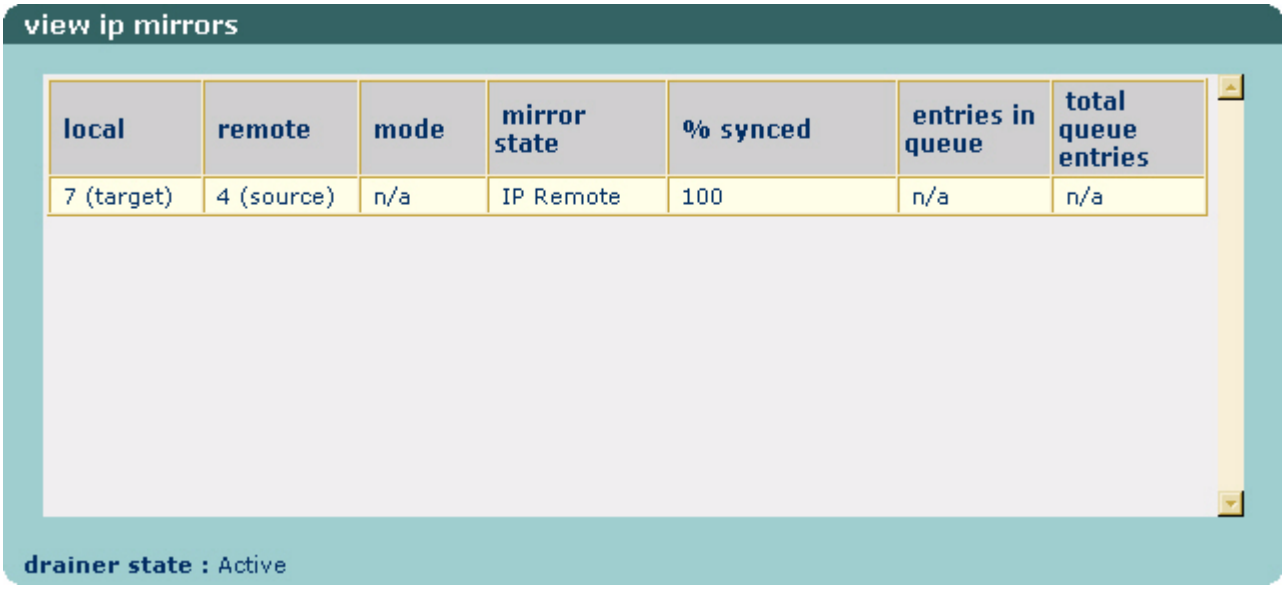

Figure 107: View ip mirrors window (from target appliance)

Table 8 describes each field in the **view ip mirrors** window.

Table 8: Fields in the view ip mirrors window

Field	Description
local	The LUN on the local appliance (the appliance you are currently using) that the mirror is using. Identifies the LUN as the mirror source or the mirror target.
remote	The LUN on the remote appliance that the mirror is using. Identifies the LUN as the mirror source or the mirror target.
mode	The type of mirroring–either synchronous or asynchronous. When viewing the mirror from the target appliance (Figure 107), the field value is **n/a**.
mirror state	The current state of the mirror. When viewing the mirror from the target appliance (Figure 107), the field value is **IP Remote**. See "About Mirror States" on page 165 for more information.
% synced	The percentage of the journal file that has been processed during a resynchronization. Refresh the window periodically to update the resynchronization's progress. The field value is 100% when either of the following occurs: ■ The resynchronization is complete. ■ You create the mirror using the **no resync** option.
entries in queue	The number of asynchronous replication requests in the queue waiting to be resynchronized. This number does not include the requests tracked in the journal file that must also be resynchronized. The field value is **n/a** for synchronous mirrors (Figure 106).
total queue entries	The total number of asynchronous replication requests in the queue (new and existing requests). The field is negative (for example, –17) when the following occurs: ■ The queue is being rebuilt in system memory. ■ The field value is the number of entries for the mirror, but there may be more. The field value is **n/a** for synchronous mirrors (Figure 106).
drainer state	The state of the AIPM queue. Field values are **Active** or **Frozen**. See "About the Drainer" on page 161 for more information.

About Mirror States

This section covers the following topics:

- Mirror States for Synchronous Mirrors
- Synchronous Mirror States and Transitions
- Mirrors States for Asynchronous Mirrors
- Asynchronous Mirror States and Transitions

You can monitor IP mirrors by checking the **mirror state** field in the **view ip mirrors** window. The mirror state only varies when you view the mirror from the source appliance. When you view the mirror from the target appliance, the mirror state is always *IP Remote*.

Mirror States for Synchronous Mirrors

Table 9 describes the synchronous IP mirror states.

Table 9: Mirror states for synchronous mirrors

Mirror State	Description
Disconnected	The network connection from the WAN to the target LUN is disabled. Once you resolve the problem and select the resynchronization type, the mirror enters one of these states: ■ Resynchronization ■ Reverse resynchronization If the `-autoresync` parameter is set to ON (see "Adjusting IP Mirror Parameters" on page 173), the mirror automatically enters one of these states.
Enabled	The state of normal synchronous IP mirroring. Replication requests are processed on the source LUN and sent to the target LUN for replication. From this state, the mirror can enter one of these states: ■ Disconnected ■ Local failed ■ Paused ■ Remote failed
Local failed	The replication request failed on the local appliance and the network connection is disabled. All existing and new replication requests are tracked in the source LUN's journal file. **Note:** The local appliance's data remains inconsistent with the remote appliance's data until the problem is resolved and the mirror is resynchronized. Once you resolve the problem and select the resynchronization type, the mirror enters one of these states: ■ Resynchronization ■ Reverse resynchronization If the `-autoresync` parameter is set to ON (see "Adjusting IP Mirror Parameters" on page 173), the mirror automatically enters one of these states.

Table 9: Mirror states for synchronous mirrors

Mirror State	Description
Paused	Mirroring is stopped. All existing and new replication requests are tracked in the source LUN's journal file. When you select a resynchronization type to resume mirroring, the mirror enters one of these states: ■ Resynchronization ■ Reverse resynchronization When resynchronization finishes, the mirror returns to the *Enabled* state.
Remote failed	The replication request failed on the remote appliance and the network connection is disabled. All existing and new replication requests are tracked in the source LUN's journal file. **Note:** The remote appliance's data remains inconsistent with the local appliance's data until the problem is resolved and the mirror is resynchronized. Once you resolve the problem and select the resynchronization type, the mirror enters one of these states: ■ Resynchronization ■ Reverse resynchronization If the `-autoresync` parameter is set to ON (see "Adjusting IP Mirror Parameters" on page 173), the mirror automatically enters one of these states.
Resynchronization	The mirror is resynchronizing by replicating data from the source LUN to the target LUN. All new replication requests are tracked in the source LUN's journal file. When resynchronization finishes, the mirror returns to the *Enabled* state.
Reverse resynchronization	The mirror is resynchronizing by replicating data from the target LUN to the source LUN. The source LUN pulls the data from the target LUN during this process. When resynchronization finishes, the mirror returns to the *Enabled* state.

Synchronous Mirror States and Transitions

Figure 108 shows the mirror states for synchronous mirrors and illustrates transitions between the states.

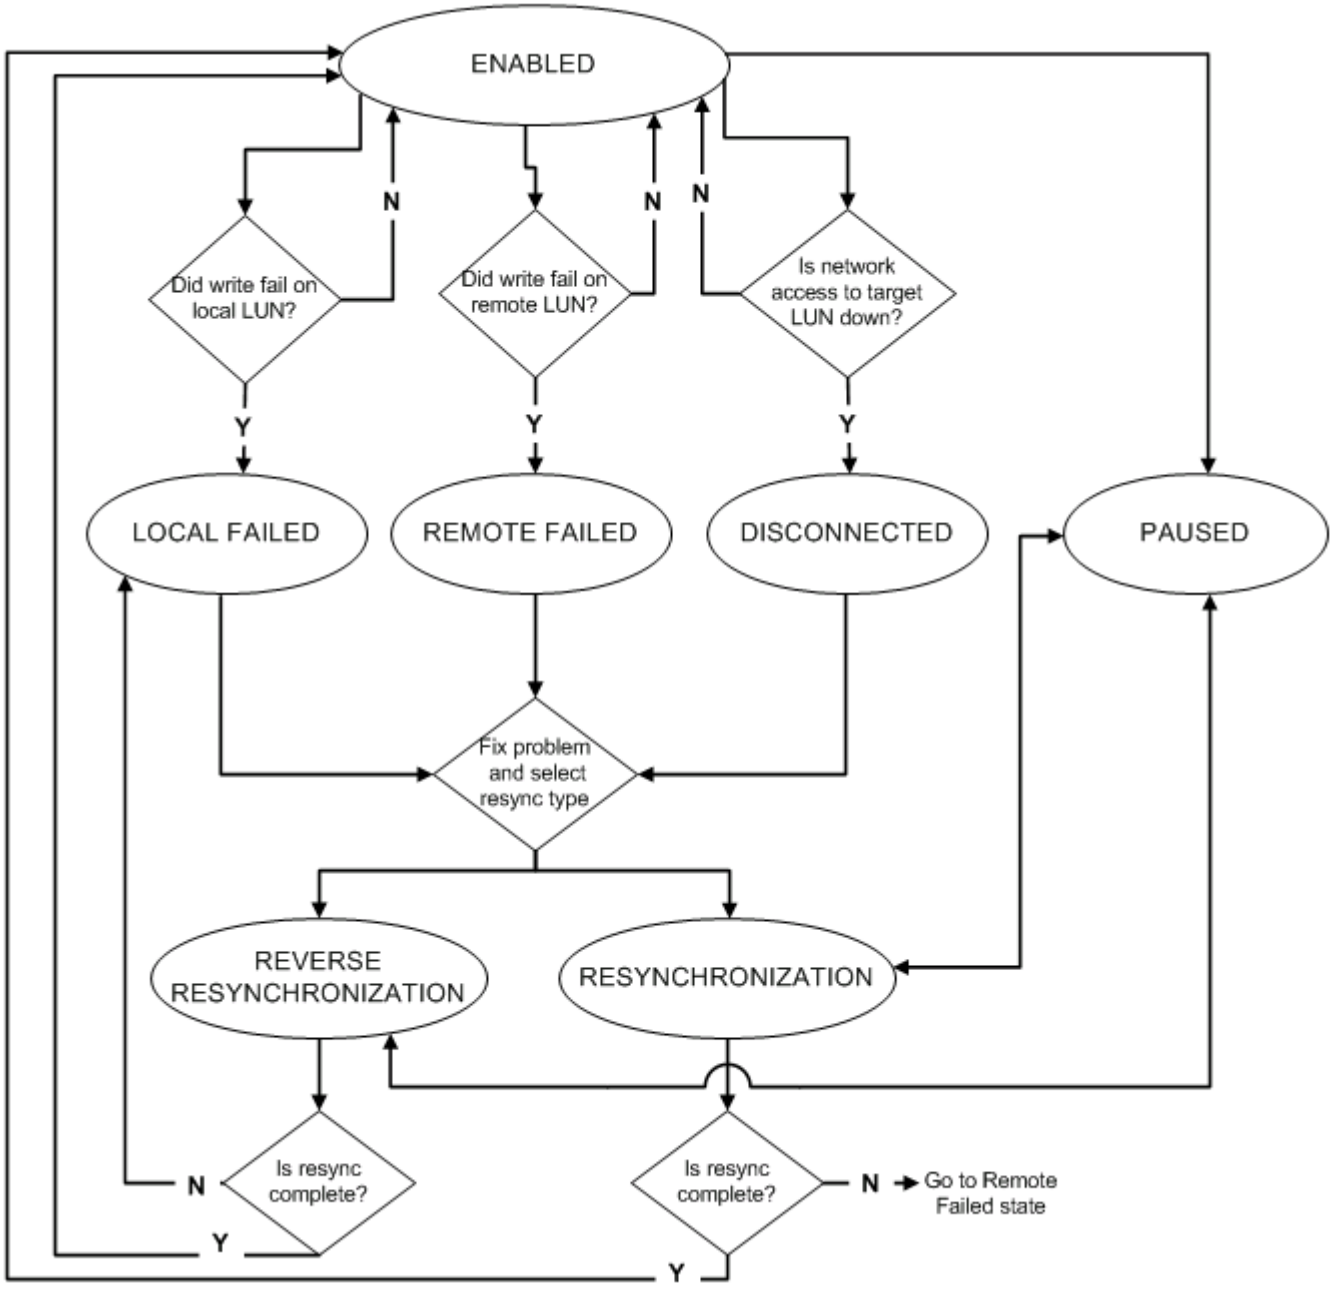

Figure 108: Synchronous mirror states and transitions

Mirrors States for Asynchronous Mirrors

Table 10 describes the asynchronous IP mirror states.

When viewing the target LUN of the mirror from the local appliance, the mirror state is always *IP Remote*.

Table 10: Mirror states for asynchronous mirrors

Mirror State	Description
Enabled	The state of normal asynchronous mirroring. Replication requests are placed in the local queue and are subsequently processed. From this state, the mirror can enter one of these states: ■ Local failed ■ Paused ■ Queue bad ■ Queue full ■ Remote dirty ■ Remote failed
Local failed	The replication request failed on the local appliance. Existing requests are drained from the queue. New requests return the following message to the host: `Sense Unit Not Ready` **Note:** The local appliance's data remains inconsistent with the remote appliance's data until the problem is resolved and the mirror is resynchronized. If you configured automatic resynchronization, the mirror enters one of the following states after you resolve the problem: ■ Resynchronization ■ Reverse resynchronization If you did not configure automatic resynchronization, the mirror enters the *Resynchronization needed* state. Choose the type of resynchronization to continue.
Paused	Mirroring is stopped. New replication requests are tracked in the source LUN's journal file while existing requests are emptied from the queue. To resume mirroring, select a resynchronization type. Depending on the type you choose, the mirror enters one of these states: ■ Resynchronization ■ Reverse resynchronization When resynchronization finishes, the mirror returns to the *Enabled* state.

Table 10: Mirror states for asynchronous mirrors

Mirror State	Description
Queue bad	A read or write failure occurs in the queue device. You must find the source of the problem to determine what type of resynchronization to perform. When resynchronization finishes, the mirror returns to the *Enabled* state.
Queue full	The AIPM queue is full. All replication requests are tracked in the source LUN's journal file. The mirror remains in this state until the following events occur: ■ The last request is emptied from the queue. ■ The source LUN and target LUN are resynchronized. If you configured automatic resynchronization, the mirror enters one of the following states once the AIPM queue has emptied: ■ Resynchronization ■ Reverse resynchronization If you did not configure automatic resynchronization, the mirror enters the *Resynchronization needed* state. Choose the type of resynchronization to continue. When resynchronization finishes, the mirror returns to the *Enabled* state.
Queue full, target needs resync	The AIPM queue is filled before resynchronization finishes. All replication requests are emptied from the queue and tracked in the source LUN's journal file. If you configured automatic resynchronization, the mirror enters one of the following states: ■ Resynchronization ■ Reverse resynchronization If you did not configure automatic resynchronization, the mirror enters the *Resynchronization needed* state. Choose the type of resynchronization to continue. When resynchronization finishes, the mirror returns to the *Enabled* state.
Remote dirty	LUNs on the remote appliance are modified while the local appliance is disabled. When the local appliance is enabled, the source LUN checks the journal file of the target LUN for modifications. Choose the type of resynchronization (resynchronization or reverse resynchronization). When resynchronization finishes, the mirror returns to the *Enabled* state.

Table 10: Mirror states for asynchronous mirrors

Mirror State	Description
Remote failed	The replication request failed on the remote appliance. All replication requests are emptied from the AIPM queue and tracked in the source LUN's journal file. **Note:** The remote appliance's data remains inconsistent with the local appliance's data until the problem is resolved and the mirror is resynchronized. If you configured automatic resynchronization, the mirror enters one of the following states after you resolve the problem: ■ Resynchronization ■ Reverse resynchronization If you did not configure automatic resynchronization, the mirror enters the *Resynchronization needed* state. Choose the type of resynchronization to continue.
Resynchronization	The mirror is resynchronized by replicating data from the source LUN to the target LUN. All new replication requests are tracked in the source LUN's journal file. When resynchronization finishes, the mirror returns to the *Enabled* state.
Resynchronization needed	The mirror enters this state if you have not configured automatic resynchronization. All replication requests are tracked in the source LUN's journal file. Choose the type of resynchronization (resynchronization or reverse resynchronization) to continue.
Reverse resynchronization	The mirror is resynchronized by replicating data from the target LUN to the source LUN. The source LUN pulls the data from the target LUN. When reverse resynchronization finishes, the mirror returns to the *Enabled* state. If a host attempts to write to the local LUN while a reverse resynchronization is in progress, the following message is returned to the host: `Sense Unit Not Ready`

Asynchronous Mirror States and Transitions

Figure 109 shows the mirror states for asynchronous mirrors and illustrates transitions between the states.

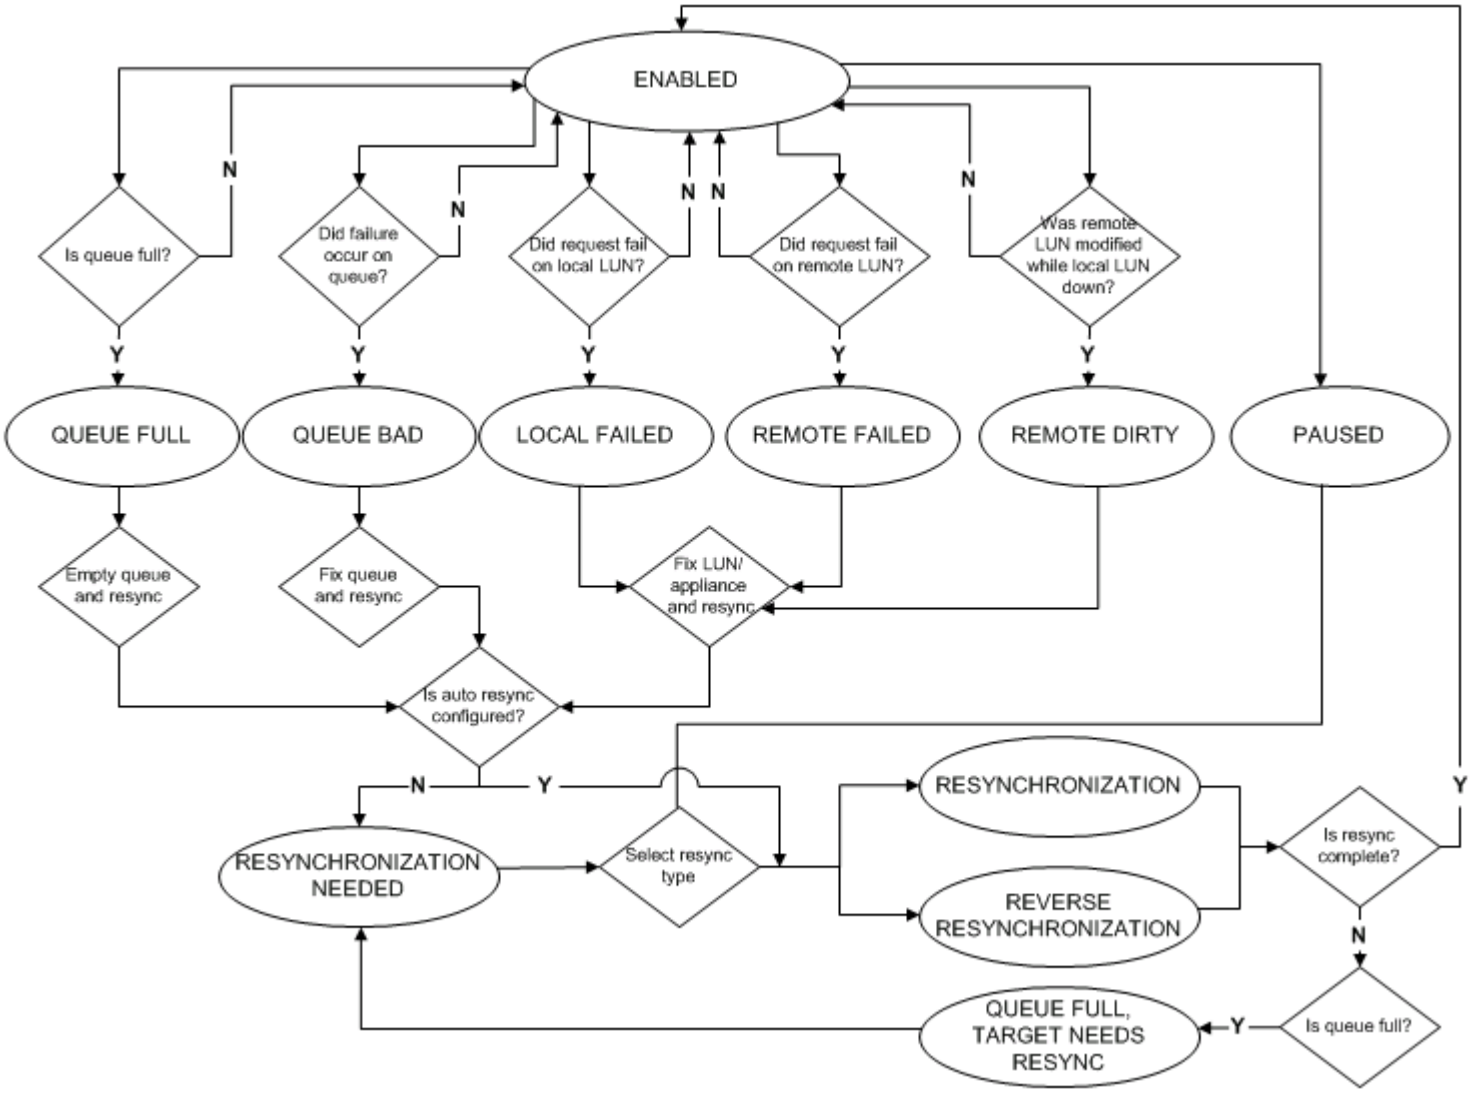

Figure 109: Asynchronous mirror states and transitions

Adjusting IP Mirror Parameters

You can use the mod param command to adjust the performance of your IP mirrors. Refer to "Mirroring Commands" in the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide* for more information about using this command.

Pausing IP Mirrors

Pausing a mirror temporarily stops data replication between the source LUN and the target LUN. Pausing is typically a manual process, but the mirror may be paused automatically if there is a communication problem with either the source LUN or target LUN. (See "About Mirror States" on page 165 for more information.)

While the mirror is paused, the host can still access the source LUN to read data. Synchronous and asynchronous mirrors handle new writes (or replication requests) differently.

This section covers the following topics:

- Synchronous Mirrors
- Asynchronous Mirrors

Synchronous Mirrors

When you pause a synchronous mirror, new replication requests are tracked in the source LUN's journal file.

For example, you pause a synchronous mirror to create a backup of the target LUN's data at a specific point in time. Once you back up the data, you resume mirroring. See "Resynchronizing IP Mirrors" on page 177 for more information.

Asynchronous Mirrors

When you pause an asynchronous mirror, new replication requests are tracked in the source LUN's journal file. Existing replication requests are drained from the AIPM queue and sent to the target LUN for completion. Unlike pausing an asynchronous mirror, freezing the drainer still allows new replication requests to be sent to the AIPM queue. See "About the Drainer" on page 161 for more information.

If you pause an asynchronous mirror to create a backup of the target LUN's data at a specific point in time, you must wait for the AIPM queue to drain completely. All existing requests must be replicated on the target LUN or the backup copy will be inaccurate.

To ensure the AIPM queue is empty, check that the value of the **total queue entries** field in the **view ip mirrors** window for the appropriate mirror is 0. If the value is not 0, click **Refresh** in your web browser to update the window. Once the value of the **total queue entries** field is 0, create the backup and then resume mirroring.

To pause an IP mirror:

1. Choose **mirrors > ip mirror > pause-resync**.

   The **pause resync ip mirrors** window opens. Figure 110 shows the mirror from the source appliance; Figure 111 shows the mirror from the target appliance.

2. Choose the mirror you want to pause.

   The **selected mirror** window opens (Figure 112).

3. Choose **pause** and click **submit**.

   The **pause resync ip mirrors** window opens, showing the mirror state is Paused (Figure 112).

You can only pause a mirror from the source appliance (the source LUN is connected to the local appliance). When viewing the mirror from the target appliance (the target LUN is connected to the local appliance), the mirror is marked with an “x” (Figure 111).

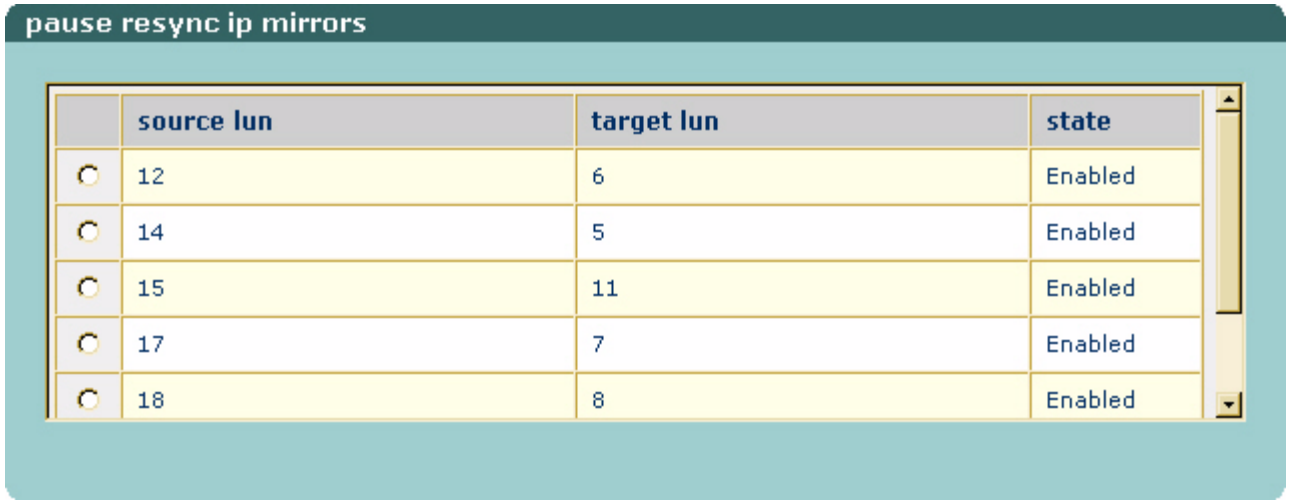

Figure 110: Pause resync ip mirrors window (from source)

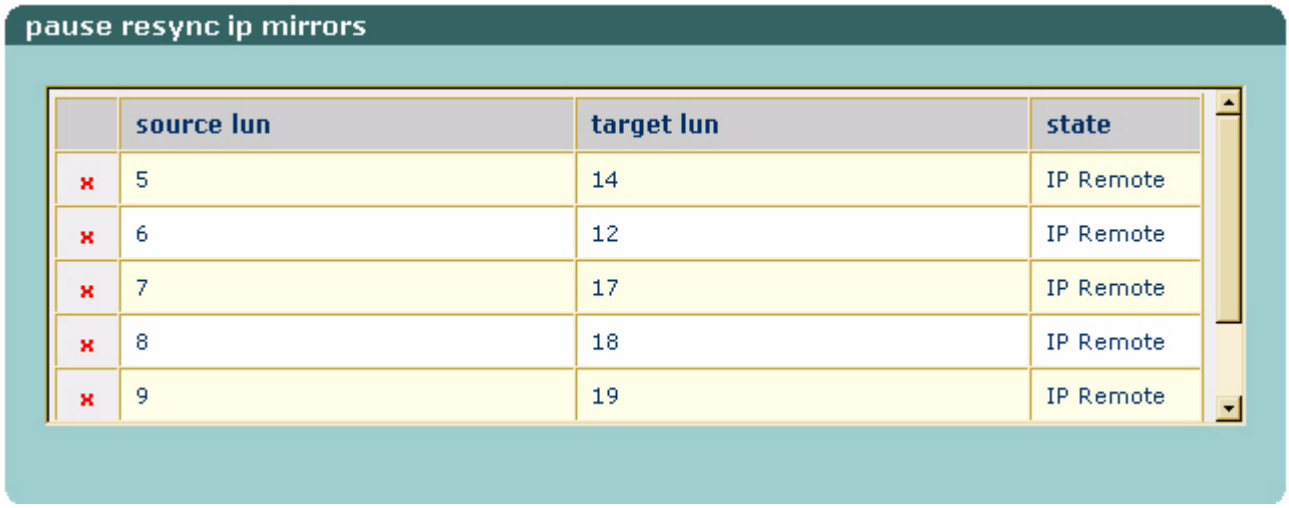

Figure 111: Pause resync ip mirrors window (from target)

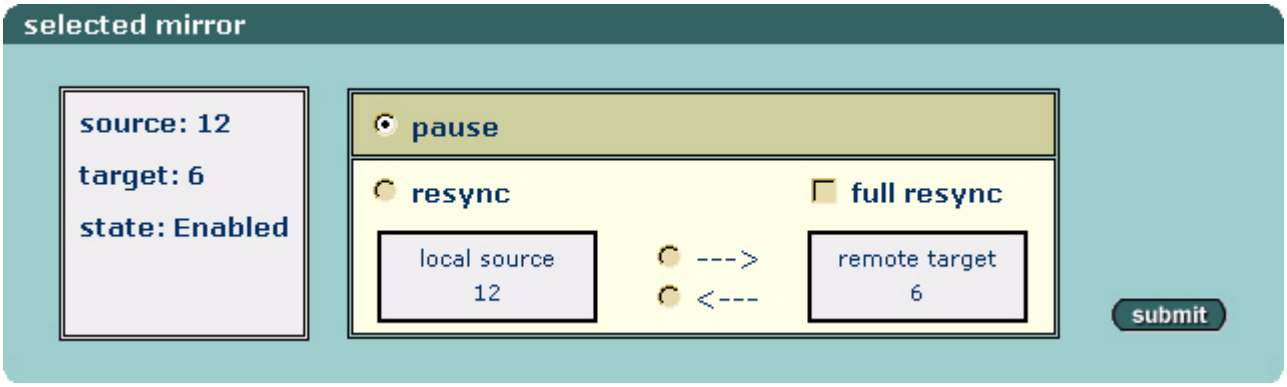

Figure 112: Selected mirror window

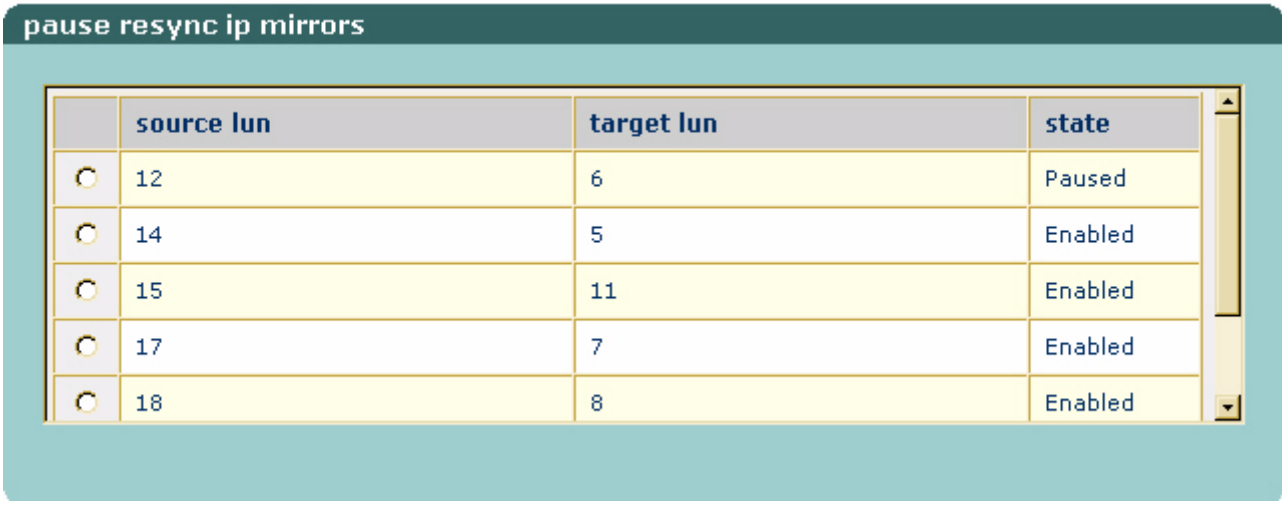

Figure 113: Pause resync ip mirrors window (paused mirror)

Resynchronizing IP Mirrors

This section covers the following topics:

- Choosing the Resynchronization Direction
- Identifying the Resynchronization Node
- Performing a Mirror Resynchronization

Mirroring can stop because:

- A failure occurred on either the local or remote appliance. See "About Mirror States" on page 165.
- You paused the mirror. See "Pausing IP Mirrors" on page 174.

Before you resume mirroring, you must resynchronize the data on the source and target LUNs so that both LUNs contain the same data.

You must choose:

- **Data to be resynchronized**—Choose **resync** to copy data received since the mirror was paused. Choose **full resync** to copy all data on the source LUN.
- **Direction of the resynchronization**—Choose forward resynchronization (source to target) to update the target LUN with the source LUN's data. Choose reverse resynchronization (target to source) to update the source LUN with the target LUN's data. In either direction, the source LUN does the work–it pushes the data (forward resynchronization) or pulls the data (reverse resynchronization).

Choosing the Resynchronization Direction

Determine the direction of the resynchronization by examining what caused mirroring to stop.

For example, you paused a mirror to create a backup of the target LUN's data. You saved the data and want to resume mirroring. You have a copy of the target LUN's data; now you can write new data to the target LUN. You choose a forward resynchronization (source to target).

In another example, you are replicating data from LUN 1 on Appliance A to LUN 4 on Appliance B (Figure 114).

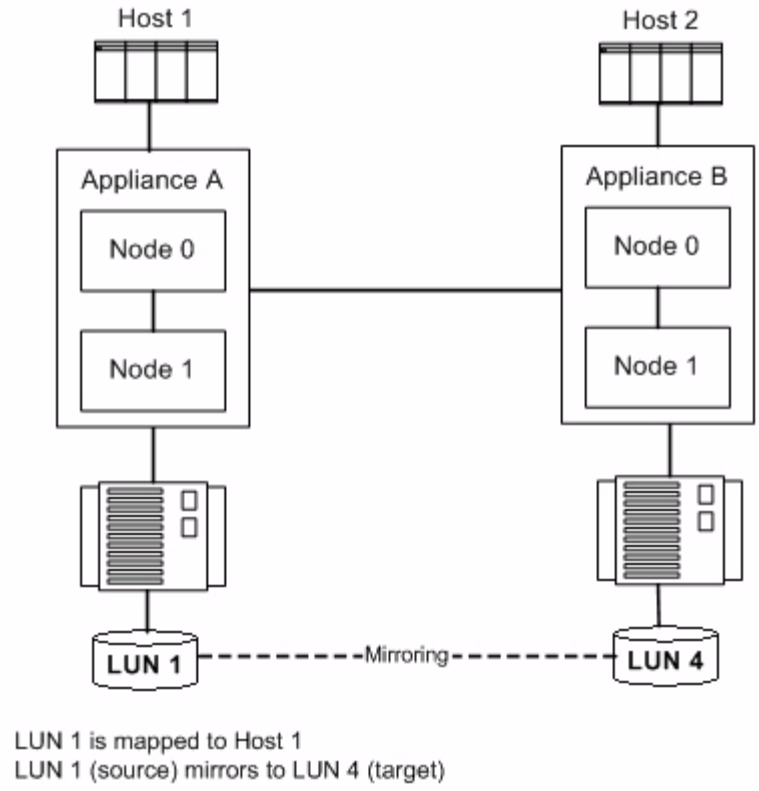

Figure 114: LUN 1 mirrors to LUN 4

Appliance A is disabled. You map Host 2 to LUN 4 to ensure read and write access continues. while Appliance A is restored. New writes are tracked in the journal file for LUN 4 (Figure 115).

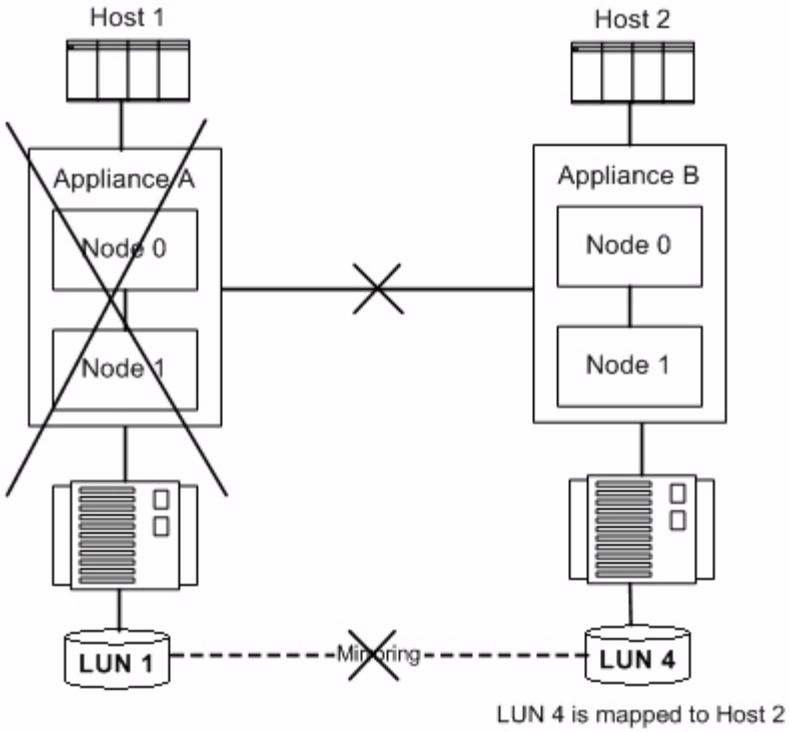

Figure 115: LUN 4 takes over while LUN 1 is disabled

When Appliance A is restored, the data on the source LUN and target LUN (LUN 1 and LUN 4) will be different.You must decide how you want to resynchronize the mirror–from the source to the target or from the target to the source. If you choose forward resynchronization (source to target), new data written to the target LUN (LUN 4) while the source LUN was down is overwritten with data from the source LUN.

If you choose reverse resynchronization (target to source), new data written to LUN 4 is added to LUN 1. The source LUN (LUN 1) checks the journal file of the target LUN (LUN 4) and copies the data from the target LUN.

In either direction, the source LUN does the work–it pushes the data (forward resynchronization) or pulls the data (reverse resynchronization).

Identifying the Resynchronization Node

When resynchronizing an IP mirror, the appliance software designates one appliance node as the resynchronization node. You restart and resynchronize a paused mirror from the resynchronization node. The system prompts you if you attempt to start resynchronization from the peer node. You must switch to the appropriate node to complete the resynchronization.

If the network connection between the peer nodes is disabled when you attempt resynchronization, the resynchronization fails. The resynchronization node cannot read the peer node's journal file because the peer node has the journal file open. Fix the network connection or shut down the peer node so the resynchronization node can retrieve data from the peer node's journal file.

Performing a Mirror Resynchronization

To resynchronize an IP mirror:

1. Choose **mirrors > ip mirror > pause-resync**.

   The **pause resync ip mirrors** window opens.

2. Choose the mirror you want to resynchronize.

   The **selected mirror** window opens Figure 116).

3. Choose **resync** or **full resync**.

4. Choose the resynchronization direction–from source to target (forward) or from target to source (reverse).

5. Click **submit**.

   The **pause resync ip mirrors** window opens, showing the mirror state is Resync (Figure 117).

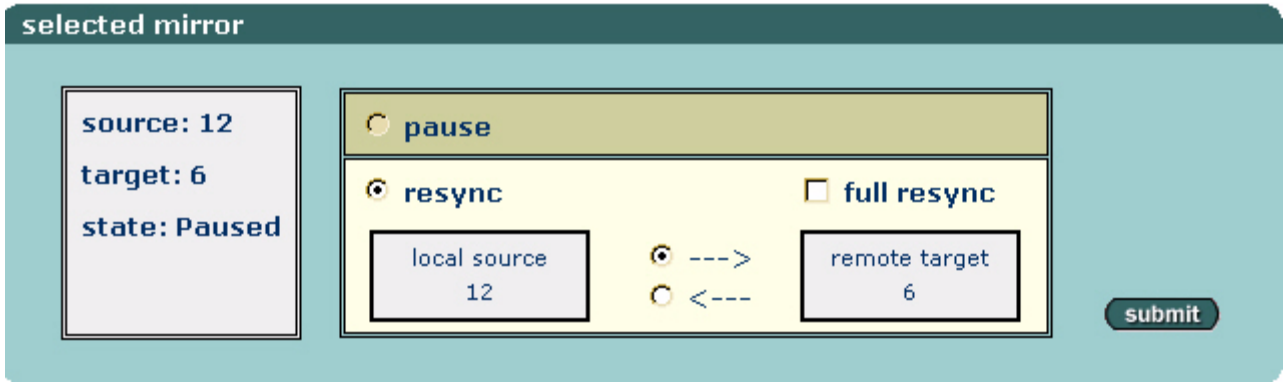

Figure 116: Selected mirror window

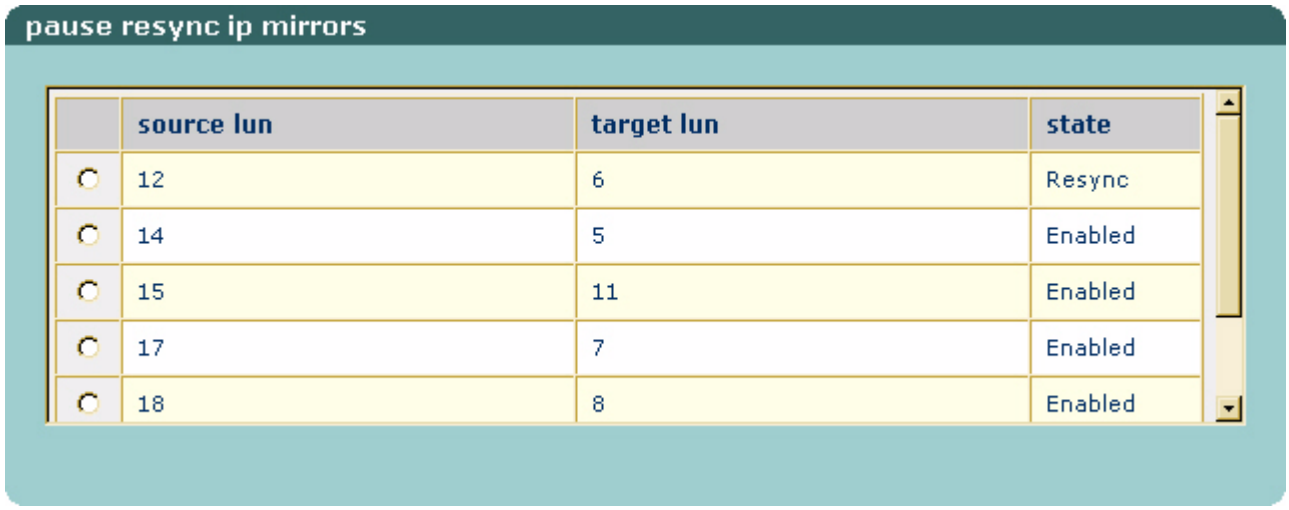

Figure 117: Pause resync ip mirrors window (resynchronizing mirror)

Deleting IP Mirrors

You may have to delete an IP mirror if:

- You have scheduled maintenance on the storage connected to the local appliance or the remote appliance.
- You want to do a full resynchronization of the source and target LUNs.
- You need to remove the mirror node. You must delete all existing mirrors first. See "Removing Mirror Nodes" on page 183.

You can delete a mirror from either the source appliance or the target appliance. If you delete a mirror from the source appliance, it is automatically deleted from the target appliance as well. If you delete a mirror from the target appliance, the mirror still exists on the source appliance. If you resynchronize that mirror on the source appliance, it is recreated on the target appliance.

To delete an IP mirror:

1. Choose **mirrors > ip mirror > break ip mirror**.

   The **delete ip mirrors** window opens (Figure 118).
2. Choose the IP mirror you want to delete.
3. Click **submit**.

   The **view ip mirrors** window opens. The mirror you deleted is not included.

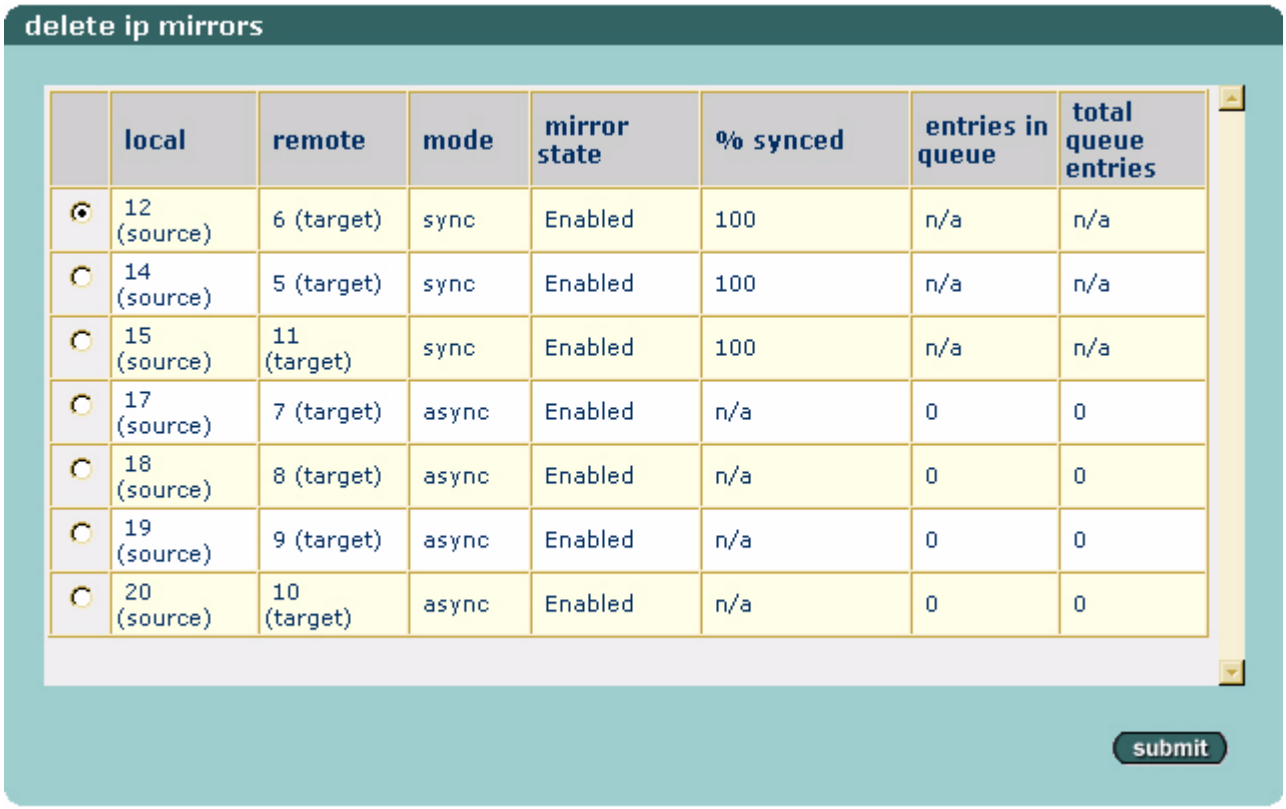

Figure 118: Delete ip mirrors window

Removing Mirror Nodes

If a failure occurs on a mirror node and you need to replace the node, you must remove the mirror relationship first. Be sure to stop all mirroring before removing the mirror node (see "Deleting IP Mirrors" on page 182). Once you have replaced the node, you can create a new mirror node relationship.

To delete a mirror node:

1. Choose **nodes > remove mirror node**.

   The **remove mirror node** window opens ((Figure 119)).

2. Click **delete**.

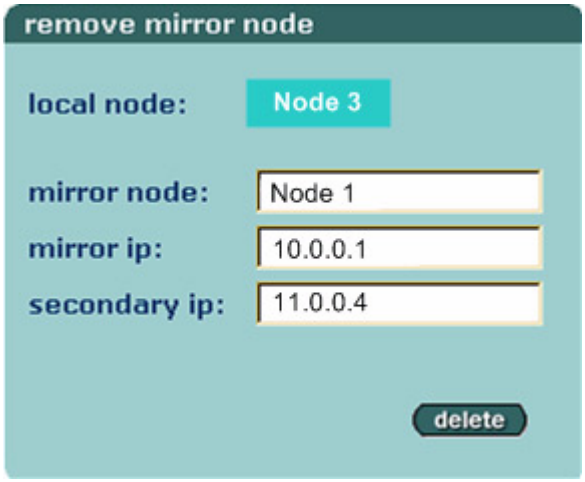

Figure 119: Remove mirror node window

Managing Point-in-Time Images

12

This chapter describes how to manage point-in-time images. It covers the following topics:

- About Point-in-Time Images, page 186
- Ensuring Data Accuracy, page 188
- Creating a Point-in-Time Image Target, page 189
- Creating a Point-in-Time Image, page 192
- Viewing Available Point-in-Time Image Targets, page 194
- Deleting a Point-in-Time Image, page 195
- Deleting a Point-in-Time Image Target, page 196

About Point-in-Time Images

A point-in-time image is a copy of data as it exists at a specific point in time. Creating a point-in-time image is similar to saving a text document periodically as you are working on it. When you save a text document, you overwrite the version you saved previously; when you create point-in-time images, you create individual copies (Figure 120).

About the Backing Store

You assign a source LUN and target LUN to a point-in-time image; however, you do not copy data to the target LUN. The point-in-time image target is a virtual LUN you create from another LUN, which is the *backing store* for the target. The point-in-time image target points to the source LUN and to the backing store for the data.

The size of the backing store can be smaller than the size of the source LUN because you are not storing all of the source LUN's data on the backing store. Before new data is written to the source LUN, the current data is copied to the backing store. For example, block 10 on the source LUN is changing from A to B. Block 10 with A is copied to the backing store. Then, block 10 on the source LUN is changed to B. When reading the blocks on the target LUN, the appliance software determines which LUN has the data required—either the source LUN or the backing store.

Note: Ensure that there is enough space on the backing store to accommodate any updates received. When there is no space available on the backing store, any point-in-time image targets that you create from it are no longer valid.

Point-in-Time Image Attributes

Point-in-time images are read-only or read-write. They are useful for backups of application data, testing, or background processing of production data. For example, you mirror a database from LUN 0 to LUN 1. You can create a point-in-time image of the database copy on LUN 1 at regular intervals. If your database application fails and sends bad data to your source, it will also send bad data to your copy. You can use the point-in-time image as a backup to help restore your database.

Typically, you do not keep point-in-time images for a long period of time. If you use point-in-time images to create copies of data at regular intervals, you can delete older copies as newer copies are made.

Figure 120 shows how the point-in-time image process works. LUN 0 is the source LUN that contains the database application. LUN 1 and LUN 2 are the backing stores. From LUN 1, you create two point-in-time image targets (LUN 3 and LUN 4). From LUN 2, you create one point-in-time image target (LUN 5). You create point-in-time images of LUN 0 at 12:00 pm, 1:00 pm, and 2:00 pm and assign each image created to a target (LUN 3, LUN 4, and LUN 5).

You map LUN 5 to Host 2 to create a backup of the image created at 2:00 pm. The appliance software reads the blocks on LUN 5, which points the software to the appropriate LUN (source LUN or backing store) for the data.

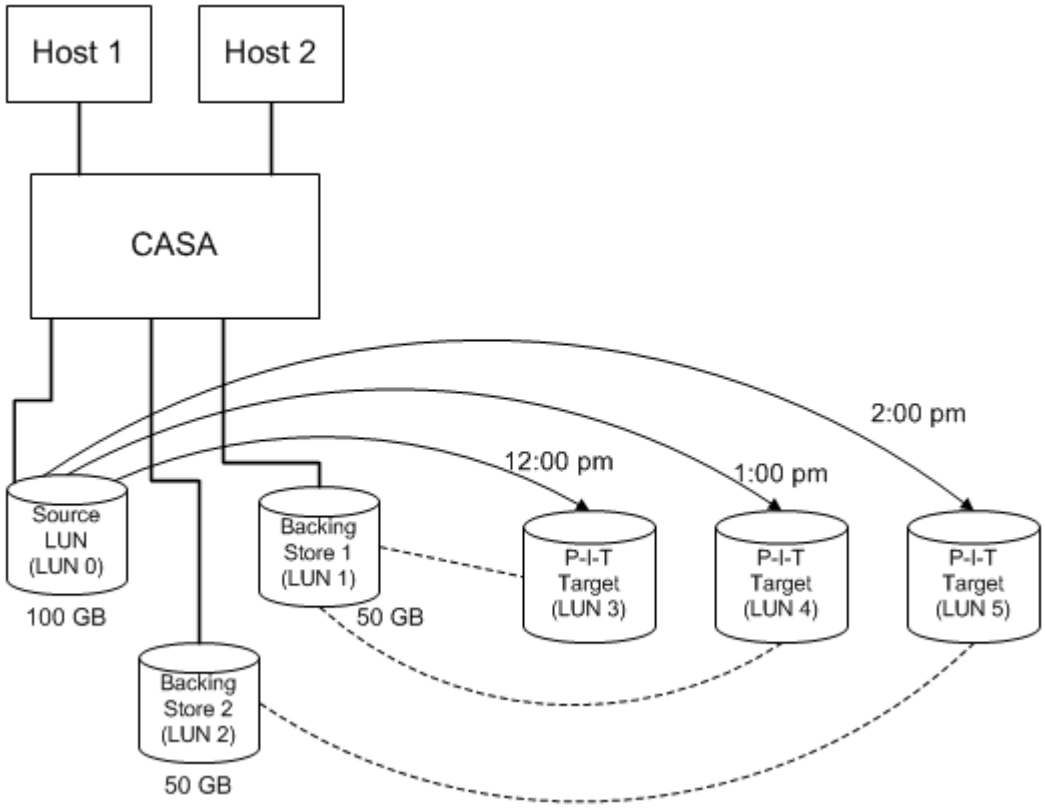

Figure 120: Point-in-time image process

Ensuring Data Accuracy

Before you create a point-in-time image, ensure that the source LUN contains the most current data. Verify that data in the application buffers has been copied to the file system and that data in the host cache has been copied (or flushed) to disk.

When you use an application on a host to store or retrieve data on a disk, the following components work behind the scenes to complete this task:

- The application stores data in blocks (called buffers) in its memory. How the application was created determines how long the data remains in the buffers. Periodically, data is moved from the buffers to a file system. The file system maps these buffers to locations on a physical disk. The file system uses large blocks of memory on the host (called cache blocks) to:
  - Store read requests and write requests in the host memory so you can quickly access frequently used files.
  - Bundle write requests to move data to the file system more efficiently.
- Host operating systems handle the caching process differently, depending on the application type, performance, and reliability you want. Therefore, if data on the disk changes when the host has data in cache, it can take time for the host to recognize the changes.
- Controllers and most disk drives maintain caches of varying sizes and policies that may affect data even after it has been removed from the host cache. Usually these caches are maintained for a specific period of time in case a power failure occurs. However, if you remove a disk from a controller while data is in the cache, that data is lost.

Because there are several levels of caching that occur, it can be difficult to determine where a particular block of data is located and, therefore, take an accurate point-in-time image.

For example, if you create a point-in-time image of an application's file system (copy all files to another location), the application may still have data in the memory buffers. If you create a point-in-time image of a disk (copy all of the blocks to another disk), the host will likely have data in its cache.

Before you create a point-in-time image:

- Stop or pause the application to ensure that the contents of all application buffers are moved to the file system.
- Unmount the file system to ensure that the cache is moved to the host.

Creating a Point-in-Time Image Target

Before you create a point-in-time image, you must designate a LUN as the point-in-time image target. To create a target, you select the LUN to be the backing store for the target. When you select the backing store, the appliance software creates the target and assigns a number to it.

You can create multiple point-in-time image targets from the LUN you use as the backing store. You use each target LUN to create a point-in-time image. This is useful if you want to designate a specific physical LUN for all point-in-time images.

To create a point-in-time image target:

1. Choose **point-in-time image > create targets**.

   The **create pt image targets** window opens (Figure 121).

2. Choose the LUN to be the backing store in the **select lun** field.

   Note: To view details about the LUN, click **show lun info**. The **LUN info** window opens (Figure 122).

3. Click **submit** in the create pt image targets window.

   The **view pt image targets** window opens (Figure 123). The appliance software assigns a number to the point-in-time image target based on the next available LUN. For example, the last available LUN is 2 (Figure 121). When you create a target, the appliance software assigns LUN 3 to it (Figure 123).

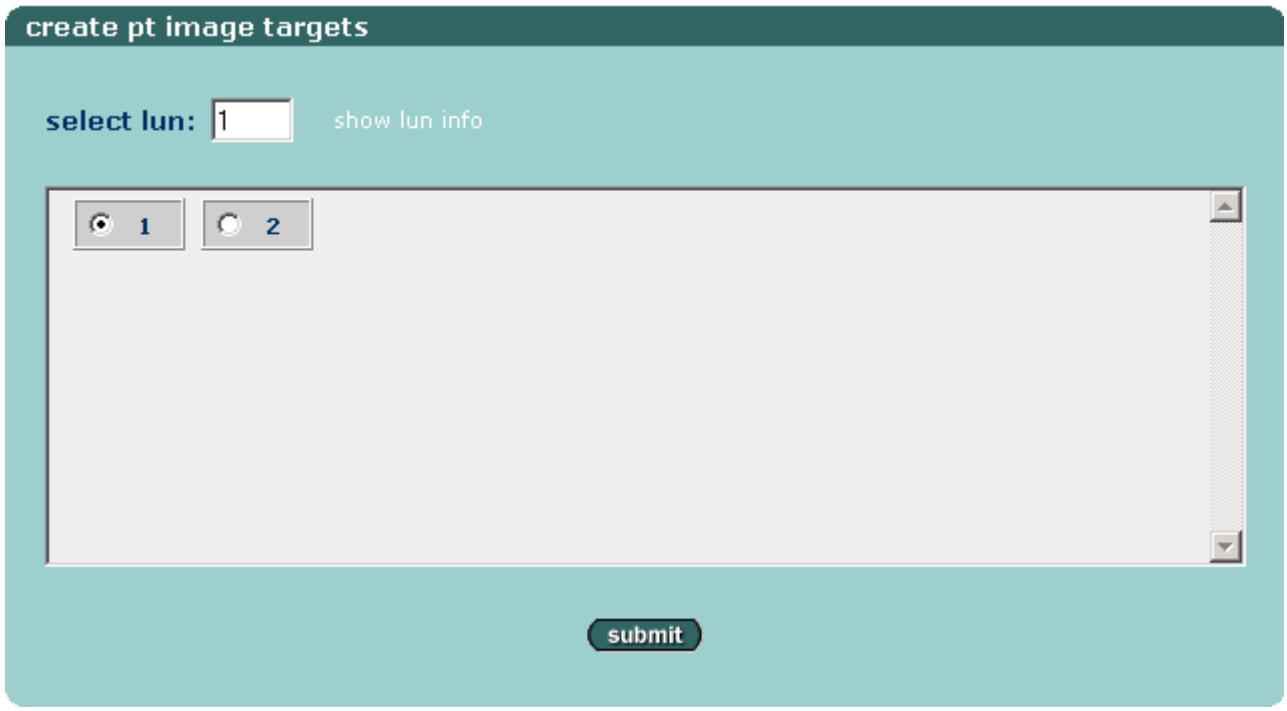

Figure 121: Create pt image targets window

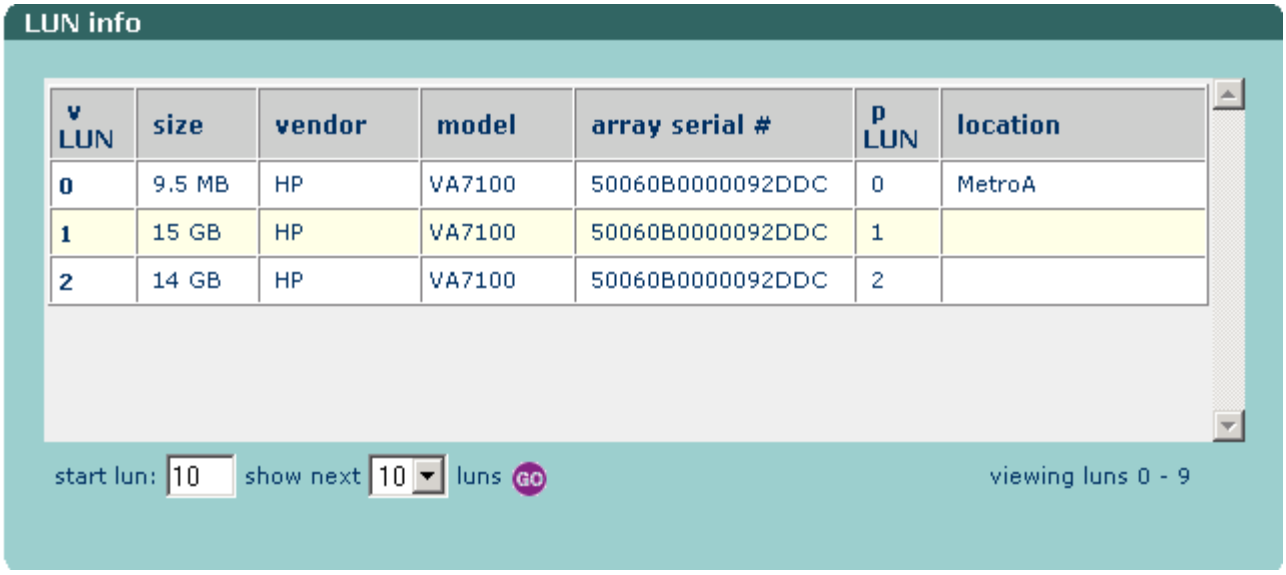

Figure 122: LUN info window

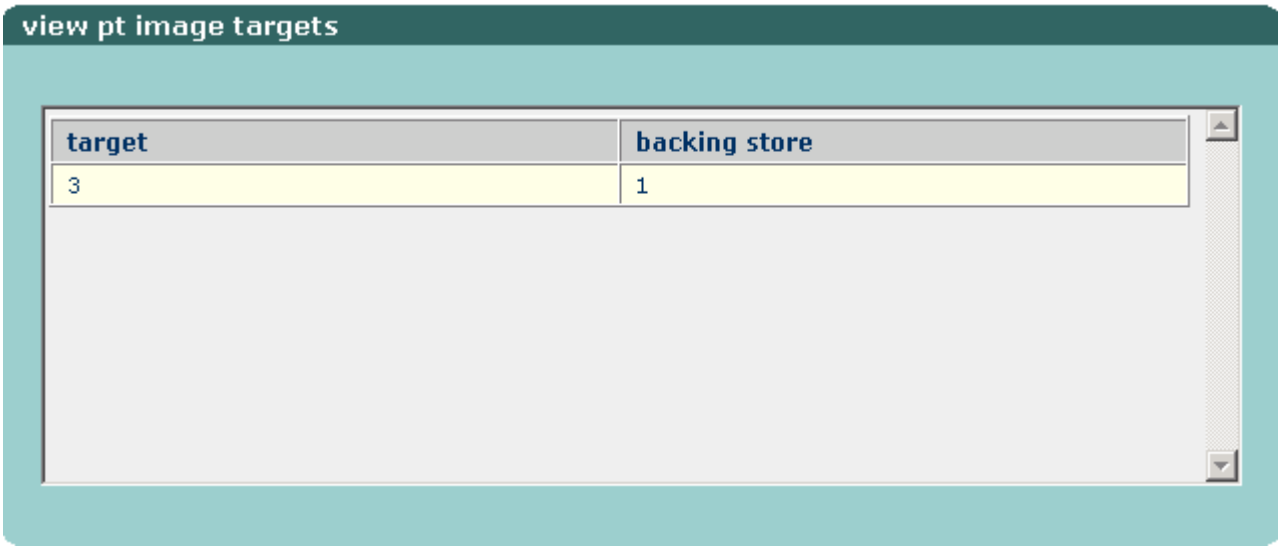

Figure 123: View pt image targets window (displaying a single target)

Figure 124 shows two point-in-time image targets created from LUN 1 and one point-in-time image target created from LUN 2.

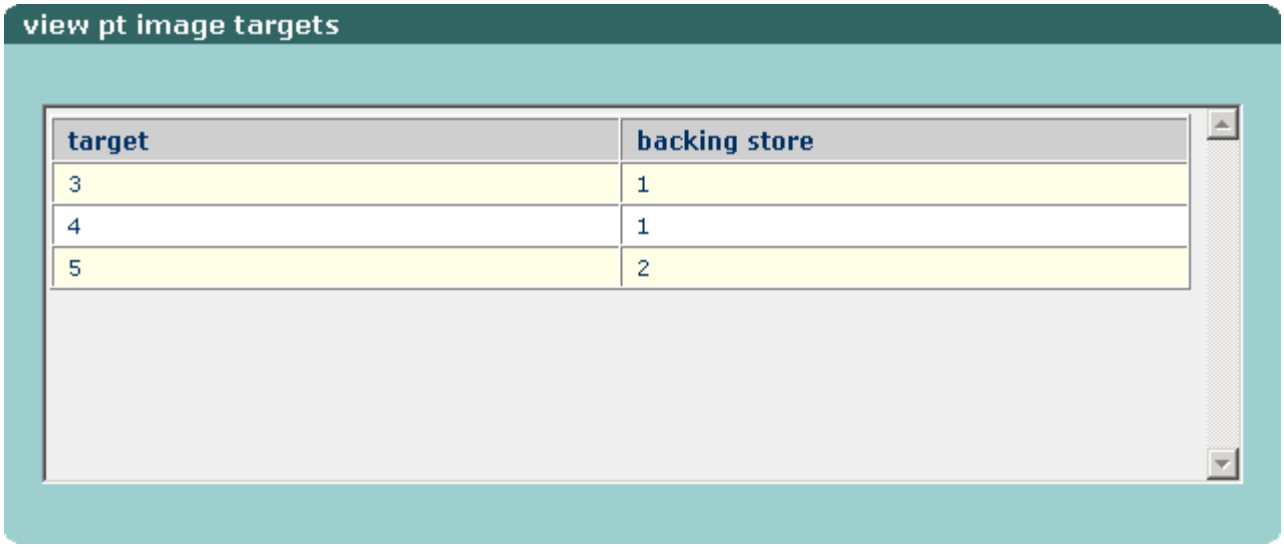

Figure 124: View pt image targets window (displaying multiple targets)

Creating a Point-in-Time Image

After you create a target, you create a point-in-time image.

To create a point-in-time image:

1. Choose **point-in-time image > create point in time image**.

   The **create point in time image** window opens (Figure 125).
2. Choose the source LUN in the **select source lun** field.
3. Choose the target LUN in the **select target lun** field. The only target LUNs available are those you created as point-in-time image targets.
4. Choose the size of the image in the **block size** field. The default size is 8K, although you can choose from 4K to 128K.
5. Click **submit**.

   The **view point in time images** window opens, displaying the point-in-time image you created (Figure 126).

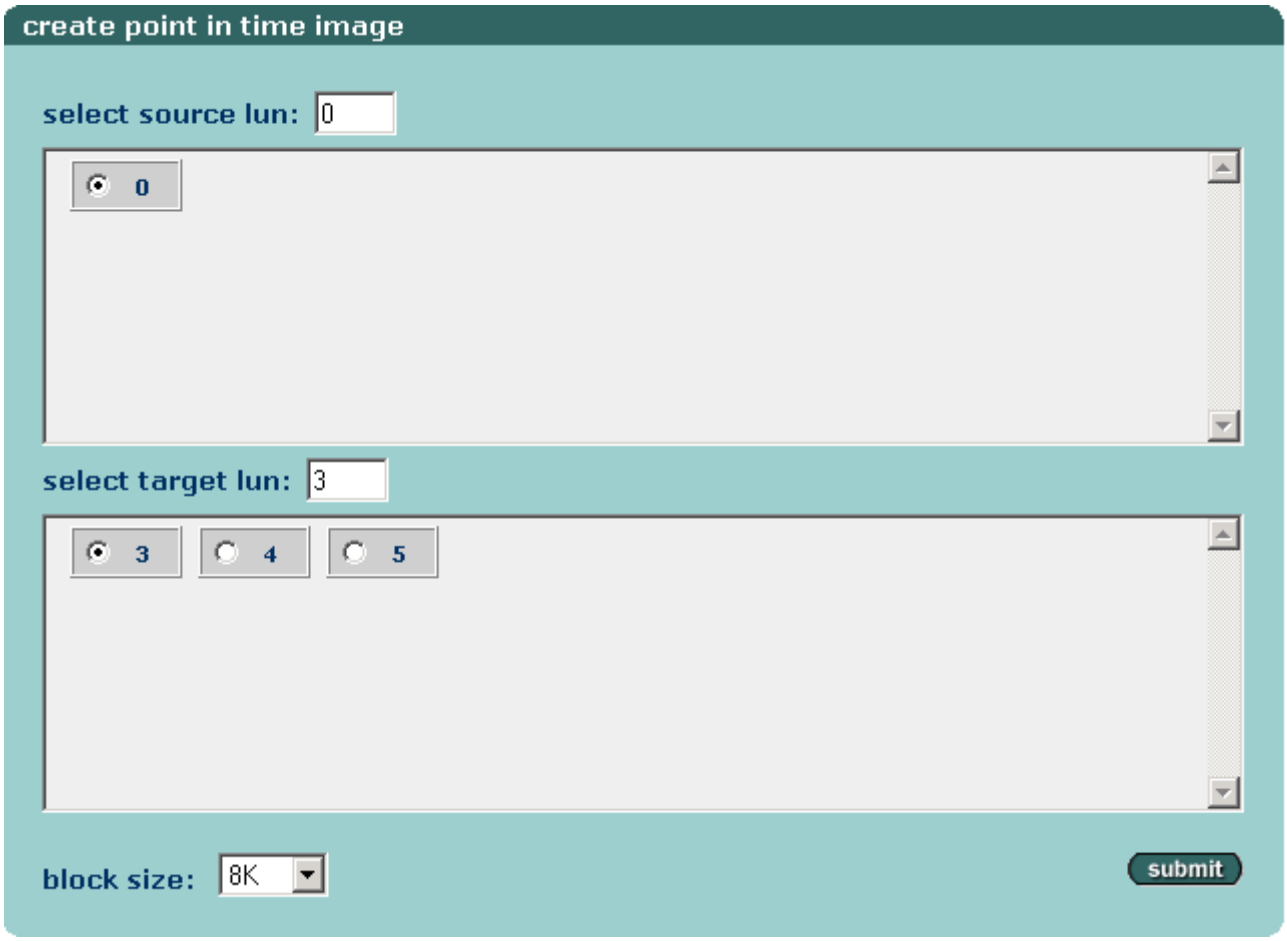

Figure 125: Create point in time image window

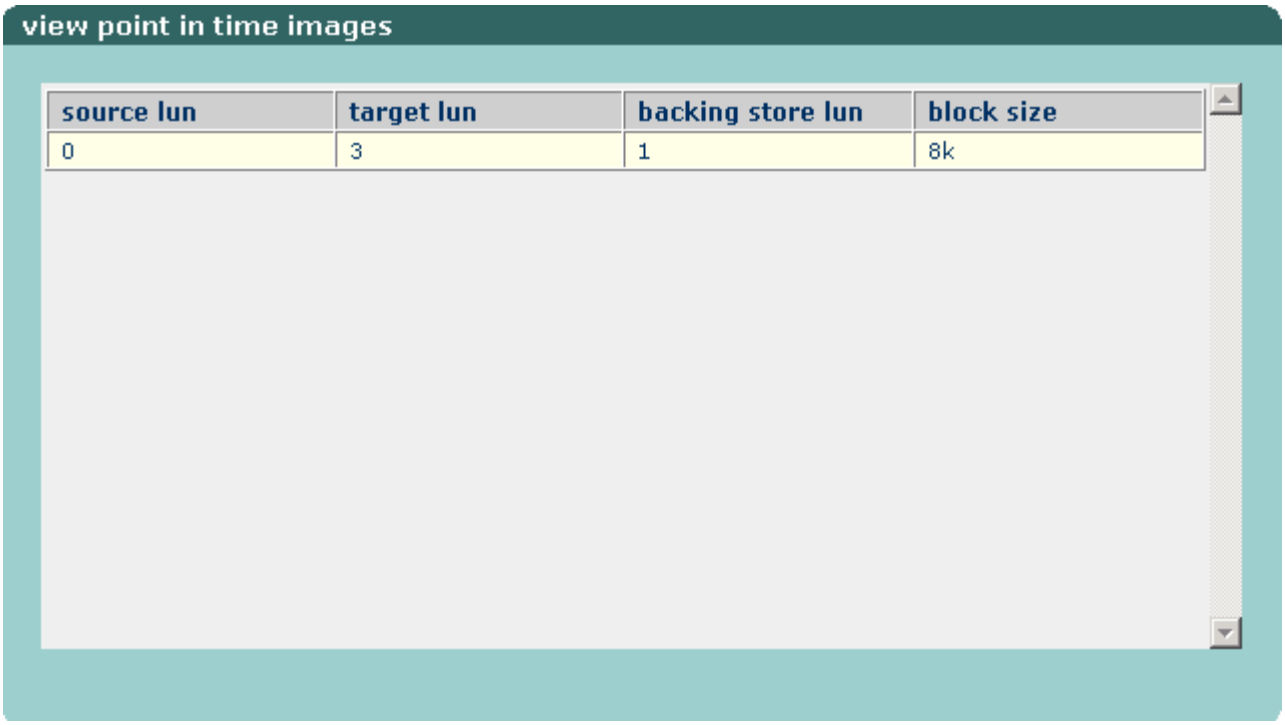

Figure 126: View point in time images window

Viewing Available Point-in-Time Image Targets

When you create a point-in-time image, you cannot use the same target LUN to create another point-in-time image. The target becomes available only when you delete the point-in-time image. You can then use that target to create another image.

To view available point-in-time image targets:

1. Choose **point-in-time image > available targets**.

   The **view pt image targets** window opens (Figure 127).

Because you used LUN 3 to create a point-in-time image (Figure 125), LUN 3 is not listed as an available target.

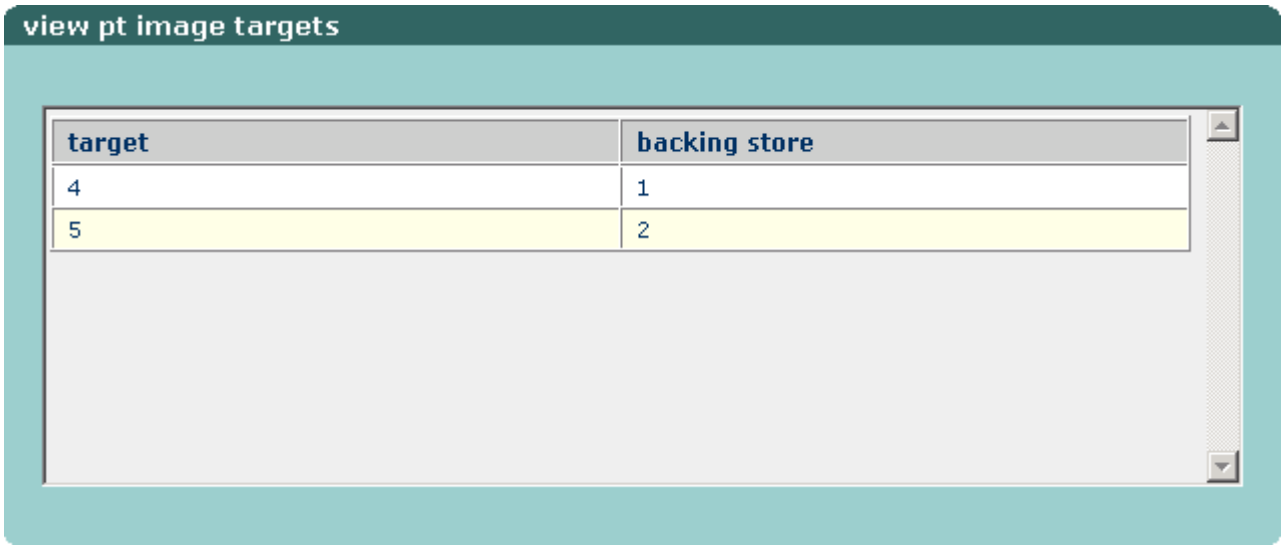

Figure 127: View pt image targets window

Deleting a Point-in-Time Image

You can delete a point-in-time image that you no longer need. You must delete the point-in-time image before you delete the target LUN for the image.

To delete a point-in-time image:

1. Choose **point-in-time image > break pt image**.

   The **delete point in time image** window opens (Figure 128).
2. Choose the point-in-time image you want to delete.
3. Click **submit**.

   The **view point in time images** window opens. The image you deleted is not listed.

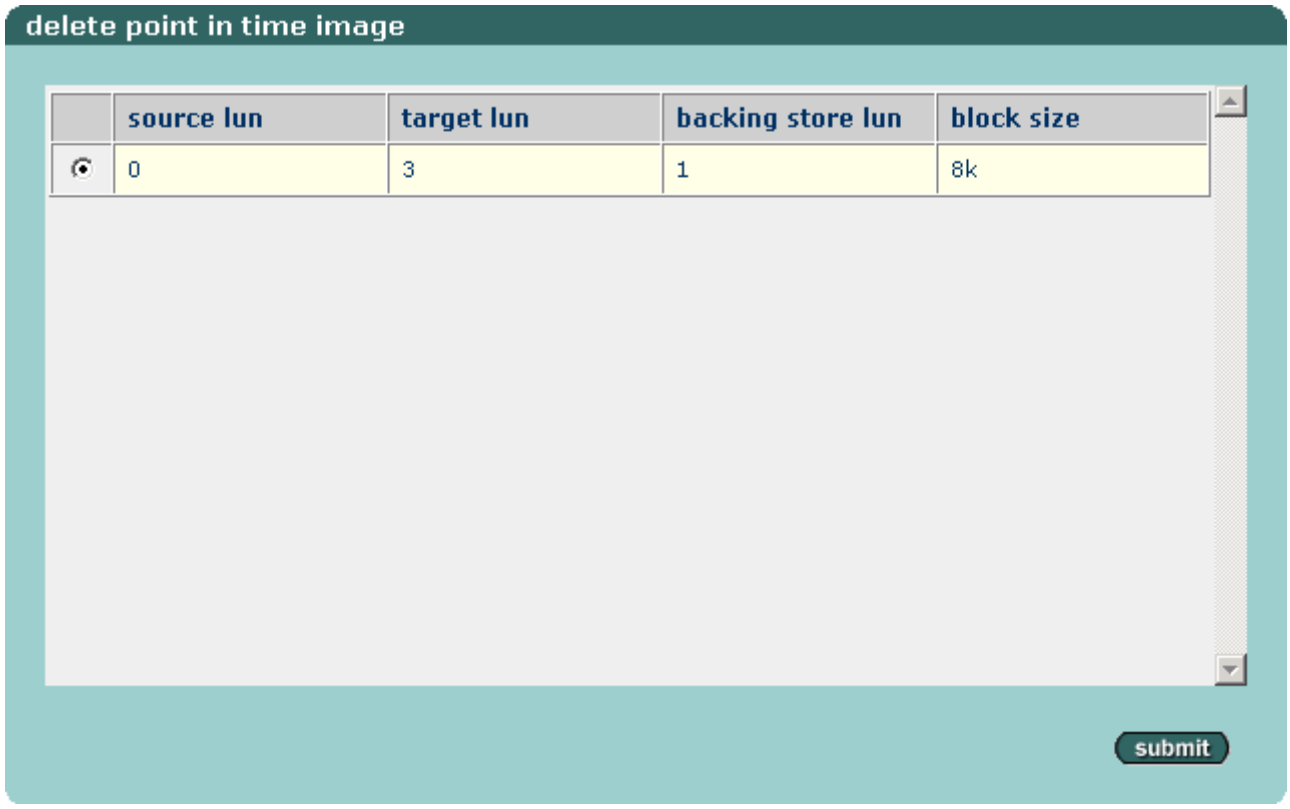

Figure 128: Delete point in time image window

Deleting a Point-in-Time Image Target

Note: Before you delete a point-in-time image target, delete the associated point-in-time image.

To delete a point-in-time image target:

1. Choose **point-in-time image > delete targets**.
   The **delete pt image target** window opens (Figure 129).
2. Choose the target you want to delete.
3. Click **delete**.
   The **view pt image targets** window opens. The target you deleted is not listed.

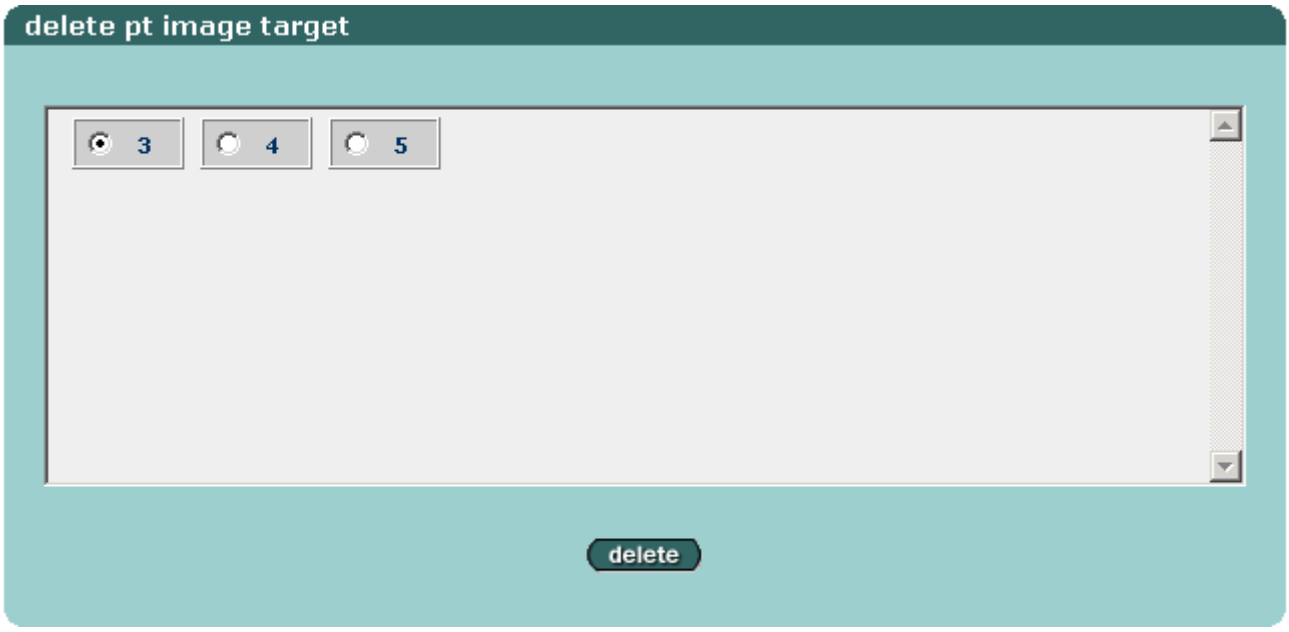

Figure 129: Delete pt image target window

Using the Switch Interface

13

You can use the switch command line interface (CLI), called **switchadmin**, to manage Brocade switches that you have connected to the appliance. You use a Telnet connection to log in to and send commands to the switch.

Note: The switch CLI is supported on Brocade switches only.

This chapter covers the following topics:

- Registering Switches, page 198
- Managing Multiple Switch Connections, page 199
- Command Types, page 200

Registering Switches

Before you can log in to a switch using the switch CLI, you must register the switch in the authorization file.

If you attempt to connect to a switch without authorizing it first, the following error message appears:

```
switchadmin: connect xx.x.xxx.xxx user password
ERROR:
IP address not authorized.
```

To register a Brocade switch:

1. Create a text file using a text editor.
2. Enter each switch IP address on a separate line:

   ```
   15.200.30.20
   15.200.30.21
   15.201.31.20
   ```

3. Save the file to the following directory:

   ```
   C:\sanlink\config\security\brocade.txt
   ```

4. You must duplicate the authorization file on the peer node. If you add, change, or delete switches in the future, you do not have to add any software services on the nodes.

Managing Multiple Switch Connections

You can use the switch CLI to manage multiple switch connections concurrently.

When you log in to the user interface from your web browser, the user interface assigns a ticket to this browser connection. When you use the switch CLI to connect to a switch, the ticket identifies the switch connection, which is called a session.

You cannot use the same session to connect to a different switch. However, you can open another browser connection to the user interface and establish a session to a different switch, while the first browser connection and session are still open. This enables you to create and manage multiple switch connections concurrently.

Note: The switch CLI supports multiple switch connections to the same switch, but the current Brocade switch software does not.

For more information about tickets, refer to "Defining Specific User Logins" in the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide*.

Command Types

There are three command types:

- SMS Commands
- Brocade Switch Commands
- Pseudo commands (page scrolling)

SMS Commands

There are two supported SMS (CASA Management Server) commands:

- Connect
- Logout

Connect

To connect to a switch:

1. Choose **utilities > switchadmin**.

   The **switch command line interface** window opens (Figure 130).
2. In the command line field (**A**), enter:

   ```
   connect [switch ip address] [username] [password]
   ```

Logout

To log out of the switch and terminate the telnet connection:

1. Choose **utilites > switchadmin**.

   The **switchadmin command line interface** window opens.
2. In the command line field, enter:

   ```
   logout
   ```

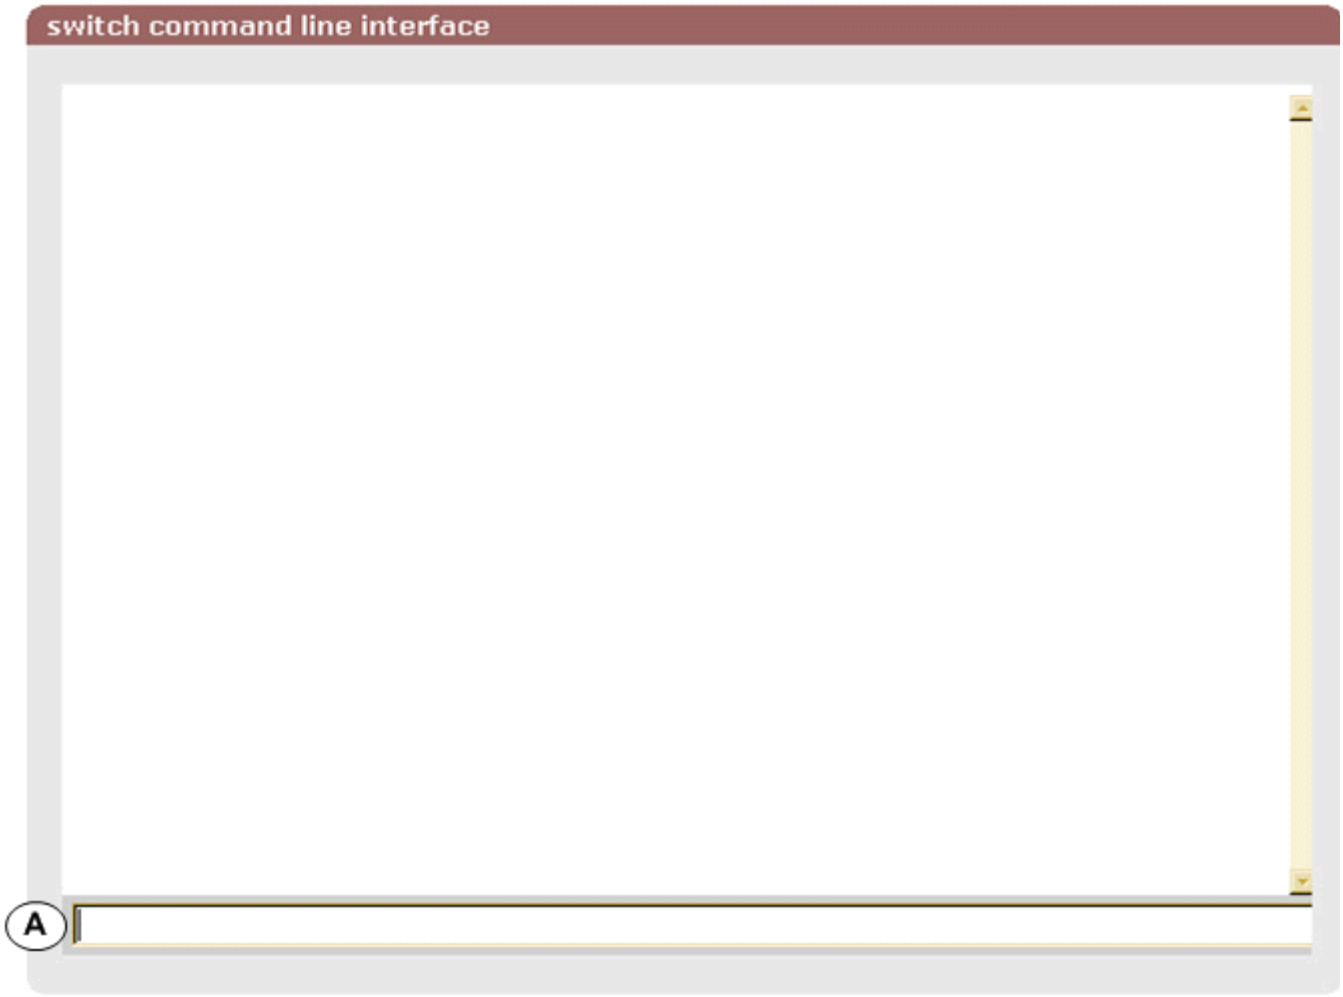

Figure 130: Switch command line interface window

Brocade Switch Commands

Table 11 lists the Brocade switch commands. Refer to your Brocade switch documentation for more information about using each command.

Table 11: Brocade Switch Commands

- aliadd
- alicreate
- alidelete
- aliremove
- alishow
- cfgadd
- cfgclear
- cfgcreate
- cfgdelete
- cfgdisable
- cfgenable
- cfgremove
- cfgsave
- cfgtransabort
- cfgshow
- configshow
- errdump
- errshow
- fabricshow
- fanshow
- fashow
- fastatsshow
- fazoneadd
- fazonecreate
- fazonedelete
- fazoneremove
- fazoneshow
- h
- help
- gbicshow
- i
- ipaddrshow
- portcfgshow
- portdisable
- portenable
- porterrshow
- portlogclear
- portlogdump
- portlogdumpport
- portlogshow
- portlogshowport
- portperfshow
- portshow
- portstatsshow
- psshow
- qloopadd
- qloopcreate
- qloopdelete
- qloopremove
- qloopshow
- supportshow
- switchdisable
- switchenable
- switchshow
- switchstatusshow
- tempshow
- version
- zoneadd
- zonecreate
- zonedelete
- zonehelp
- zoneremove
- zoneshow

To view the syntax for a command, enter:

```
Help [command_name]
```

Note: The `supportshow` command currently takes approximately four minutes to execute; this will be remedied in a future release.

Backing Up and Restoring Files

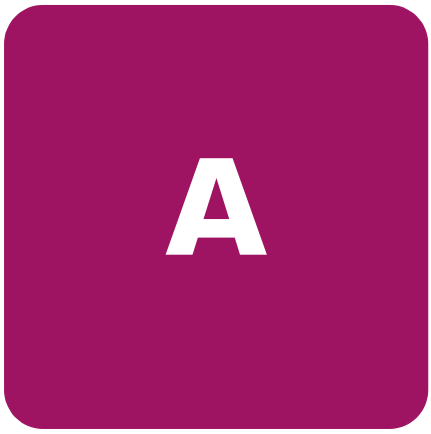

This appendix describes how to back up and restore the appliance configuration data. It covers the following topics:

- Backing Up Files, page 204
- Restoring Files, page 205

Backing Up Files

The backup utility creates a backup of the appliance configuration data and saves it on the local drive (c:\) of each node. Therefore, you need only run the backup utility on one appliance node.

To back up the appliance configuration data:

1. Choose **sanos admin > launch** in the **hp OV CASA Console** window.

   The **sanosadmin** window opens.

2. Log in to sanosadmin.

3. Enter the following command to create the **slbackup** user:

   ```
   add user -user slbackup -passwd slbackup -role guest
   ```

Note: You create the **slbackup** user once—the first time you perform a backup.

---

4. Choose **backup\restore > backup** in the **hp OV CASA Console** window.

   The backup utility window opens and the backup process begins.

5. When prompted, enter a unique comment for the backup.

   Use this comment during the restore process to identify the correct backup from which to restore. After you enter the comment, the backup utility indicates that the comment has been assigned and lists the directory in which the backup files will be stored.

   The format of the backup file name is **c:\SLBACKUP\yyMMddhhmmss**:

   - **yy**—Year
   - **MM**—Month
   - **dd**—Day of month
   - **hh**—Hour
   - **mm**—Minutes
   - **ss**—Seconds

   The **\SLBACKUP** directory has three subdirectories:

   - **ssfm**—The files saved from the shared drive
   - **node0**—The files saved from the local drive of node 0
   - **node1**—The files saved from the local drive of node1

6. When prompted to continue the backup procedure, enter **Y**.

   The backup utility copies the configuration files from the shared storage array and the local drives of each node to the backup directory on the local node. Upon completion, all configuration files are copied to the same backup directory on the peer node's local drive.

The backup utility indicates whether the backup was successful or not.

Restoring Files

You can restore appliance configuration files from one node. The restore utility restores the configuration files on both nodes.

To restore appliance configuration files from a backup:

1. Choose **backup\restore > restore** in the **hp OV CASA Console** window.

   The restore utility window opens and the restore process begins.

2. When prompted, select a restore option:

   - **F—**Restores files to the shared storage array and the local drives of both nodes. Choose this option when you are upgrading the appliance software and you want to save your current configuration.
   - **S**—Restores files to the shared storage array only. Choose this option only if you are replacing the shared storage array. Be careful when choosing this option because some configuration data is stored on the node's hard drive and some is stored on the shared storage array. As a result, the configuration data could be inconsistent.
   - **P**—Restores files to the local drive of the peer node only. Choose this option when one of the appliance nodes is down and you want to reimage that node and restore configuration data from the operational node.

3. When prompted, enter the source path for the backup files. `c:\` is the default directory. The directory **\SLBACKUP** is appended to the source path (for example, **C:\SLBACKUP**).

   The restore utility displays a list of backups, identified by:

   - **Restore number**—Index number of the entry; corresponds to the restore option selected in step 2.
   - **Backup ID**
   - **Time stamp**
   - **Comment**—Entered in step 5 of the backup procedure (see page 204)

   The backups are displayed in groups of 20. If you do not see the backup you want, press **Enter** to display the next 20 backups.

4. When prompted, enter the restore number.

   The restore utility displays the source directory, restore number, backup ID, time stamp, comment, and restore type of the backup you selected. It also lists the files to be restored.

5. Choose one of the following options to continue:

   - **Y**—Complete the restore.
   - **N**—Display the list again.
   - **Q**—Exit the utility.

6. If you selected a peer node restore in step 2, the system prompts you to log in as the administrator.

7. As each configuration file is restored, a message is displayed, listing the source and target of the copy.

The restore utility indicates whether the restore was successful or not.

Using Scripts

B

This appendix describes how to use scripts in the command line interface (CLI). It covers the following topics:

- About Scripts, page 208
- Creating a Script, page 209
- Executing a Script, page 211
- Using the Perl Module, page 212

About Scripts

You can create a script to run a series of CLI commands automatically. You can also create a script to run a series of scripts that you use frequently.

For example, you can create and run a script to perform the following tasks:

- Determine the current value of the `maxresyncthreads` parameter.
- Change the value to 50.
- Check the value to verify that the command was successful.

The script would contain the following commands:

```
get param maxresyncthreads
mod param -maxresyncthreads 50
get param maxresyncthreads
```

Creating a Script

You can create a script file that contains a ticket, which allows the script to run without user intervention. If you do not include a ticket in the script, the system will prompt you to log in before executing the script. Refer to "Defining Specific User Logins" in the *HP OpenView Continuous Access Storage Appliance Command Line Interface Reference Guide* for information about creating tickets.

The user of the ticket must have all privileges required for all commands in the script file; otherwise, certain commands may fail. See "About Roles and Privileges" on page 98 for more information about privileges.

Create the script using the editing tool of your choice. If you create the script using a text editor (such as Notepad), enter the commands as shown in the following example:

```
get param maxresyncthreads

mod param -maxresyncthreads 50

get param maxresyncthreads
```

If you create the script as a Windows batch file, enter the commands as shown in the following example:

```
sanosadmin -server 15.1.196.120 -ticket testfile get param
maxresyncthreads

sanosadmin -server 15.1.196.120 -ticket testfile mod param
-maxresyncthreads 50

sanosadmin -server 15.1.196.120 -ticket testfile get param
maxresyncthreads
```

You must specify the server on which the script will run because you do not execute batch file from sanosadmin. You must include the `-ticket` parameter for each line because the batch file can only handle one line of commands at a time. Therefore, the batch file must log in to the specified server to run each command included in the script. If the ticket file is not included, you will be prompted to log in to execute each command.

Another benefit of using a batch file is that you can run the same commands on multiple servers. For example, if you want the script to run on server 15.1.196.121, the script file would look like this:

```
sanosadmin -server 15.1.196.120 -ticket testfile get param
maxresyncthreads

sanosadmin -server 15.1.196.120 -ticket testfile mod param
-maxresyncthreads 50

sanosadmin -server 15.1.196.120 -ticket testfile get param
maxresyncthreads

sanosadmin -server 15.1.196.121 -ticket testfile get param
maxresyncthreads

sanosadmin -server 15.1.196.121 -ticket testfile mod param
-maxresyncthreads 50

sanosadmin -server 15.1.196.121 -ticket testfile get param
maxresyncthreads
```

Scripts and Comments

Comments are useful when you want other users to understand the purpose of the script and to be aware of issues. To add comments to a script, prefix comments using `rem`, `//`, or `#`. An example of including comments in a script follows:

```
get param maxresyncthreads rem check the current value of maxresyncthreads
mod param -maxresyncthreads 50 // set the value to 50
get param maxresyncthreads # see if the mod command succeeded
```

You can add a comment or a status message that users will see when the script is executed. For example, if you want to include the message "This script has completed successfully" at the end of the script, enter the text with the prefix as shown below:

```
echo This script has completed successfully.
```

Executing a Script

You can execute a script by:

- Using the `script` command in sanosadmin
- Running a batch file

To execute the script using the `script` command, log in to sanosadmin on the server on which the script will run. At the `sanosadmin` prompt, enter the name of the script and the ticket information, as shown in the following example:

```
sanosadmin -script test.txt -ticket c:\\tmp\\ticketfile
```

The command output is displayed as the script is executing.

If you want to display the commands as the script is executing, include the `set` parameter when you create the script. This is useful for debugging the script.

An example of the using the `set` parameter follows:

```
get node
get ipmirror
set on
get param maxresyncthreads
set off
mod param -maxresyncthreads 50
get param maxresyncthreads
```

When you include `set on` before the `get param maxresyncthreads` command (as shown above), the system displays the command name and command output as the script is executing. An example of the command ouptut follows:

```
+ get param maxresyncthreads
Node: LOCAL
MaxResyncThreads: 10
```

Scripts and LAN Connections

If you use more than one LAN connection when executing a script, the script may not execute. If you created the ticket using a specific IP address and you use a different IP address when you execute the `login` command, the system displays the `Invalid Ticket` message.

To avoid this error, ensure that:

- Only one LAN connection is active when you execute a script.
- LAN connections are on different subnets.

If you include a ticket in the script, ensure that the ticket has not expired or the script will fail. If you want to use script indefinitely, set the value of the `-duration` parameter for the ticket to `0` to ensure the ticket never expires.

Using the Perl Module

When you execute a script, you may need to parse the output to determine the success or failure of the commands and to locate specific information. The Perl module contains subroutines that you can use with your scripts. It communicates with the appliance using the same protocol as the CLI, but the Perl module makes the response from the appliance more accessible than the CLI output.

You can download the Perl module from the **downloads** window in the user interface:

1. Choose **utils > downloads**.

   The **downloads** window opens.
2. Choose the Perl module in the **select file to download** field.
3. Click **downloads**.
4. Choose the directory in which you want to save the Perl module.
5. Open the Perl module and access the subroutines.

glossary

This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

active/active
In a networked appliance or storage device with two nodes, an active/active controller configuration means that both data paths are being utilized for data traffic simultaneously. In a path failure scenario, all data traffic would be directed down the remaining active path. The advantage of an active/active configuration is increased bandwidth and full utilization of hardware.

active/passive
In a networked appliance or storage device with two networked nodes, an active/ passive configuration means that the active data path is being utilized for data traffic while the passive path becomes active only in an active path failure scenario. The advantage of an active/ passive configuration is system stability.

adapter
A hardware device that enables one device to interact with another. Adapters are used to connect computers with peripheral devices or networks. In a Storage Area Network (SAN) architecture, a host uses a host bus adapter (HBA) to connect to the network.

appliance
Hardware and software components bundled as a single product to manage your storage area network.

array
See disk array.

asynchronous
In IP mirroring, asynchronous transfer occurs when the host is notified as soon as the write request is dispatched, rather than waiting for an acknowledgement.

availability
The probability that a system is available at a given point in time. Averaged over time, availability expresses the expected percentage of time a system is available to perform functions.

back end
The segment of a hardware configuration controlled by a server, appliance, and/or software. From the appliance perspective, the back end includes the initiator HBA(s) up to and including the disk array(s).

back up
Make a copy of data on one storage device and save it on another storage device to prevent data loss if a system failure occurs.

backup
Copy of data that is saved on a storage device separate from the storage device on which the original data resides.

bandwidth
The transmission capacity of a link or a system, i.e., the amount of data that can be transmitted in a fixed amount of time. For digital devices, the bandwidth is expressed in bits per second (bps) or bytes per second. For analog devices, the bandwidth is expressed in cycles per second or Hertz (Hz).

bit
The smallest unit of information on a computer, a binary digit. A single bit can hold one of two values—0 or 1. You can obtain more meaningful information by combining consecutive bits into larger units. For example, a byte is composed of eight consecutive bits.

boot
To load the operating system of a computer. This action occurs automatically when a machine is powered on.

bus
The main communication path in a computer. An electrical pathway through which signals are sent from one part of the computer to another.

cascaded configuration
Collection of networking devices that are interconnected to expand the function of the device.

channel
A connection between two components.

client
A recipient of services in a client/server application.

client/server
The relationship between machines in a communications network. The client is the requesting machine, the server is the supplying machine. It is also used to describe the information management relationship between software components in a processing system.

cluster
A group of two or more computers that function together as a single entity for fault tolerance and load balancing.

clustering
Using two or more computer systems that have access to each others applications and data. Generally refers to multiple computer systems – in which each computer may be a multiprocessor system itself - that are linked together in order to handle variable workloads or to provide continued operation in the event one fails. A cluster of computer systems provides fault tolerance and/or load balancing.

data accessibility
The system requirements for network storage, such as continuous access to files and data, high performance, and protection from data loss.

data availability
The probability that a system's data is available at a given point in time. Averaged over time, data availability expresses the expected percentage of time that data is available on a system.

Data Propagation Framework
An object-oriented framework for the construction of storage applications where data flow is described as sets of objects; and data and commands are encapsulated as objects. Data follows paths, mirrored or partitioned, and changing the path changes how storage is accessed.

data reliability
Mean length of time over which data stored in a disk subsystem can be correctly retrieved, expressed in Mean Time to Data Loss (MTDL).

data transfer rate
The maximum rate at which data can be transferred from one device to another. Transfer rates are often measured in terms of megabytes per second or kilobytes per second.

disk array
An arrangement of two or more disk drives with control software that allows the suite of disks to act as a single storage subsystem. Disk arrays may be physically grouped in RAID or daisy chain fashion. The control software presents the disks' storage capacity to hosts as one or more virtual disks.

DPF
See Data Propagation Framework.

driver
The software that works to communicate between an operating system (OS) and a peripheral.

Ethernet
A protocol that utilizes CSMA/CD. Standard Ethernet runs at 10 million bits per second (10 Mbps).

expansion
The concatenation of available storage space, spread across multiple partitions into a single storage pool accessible as a single LUN.

fabric
Any system or structure consisting of connected parts.

fabric switch
A type of switch in which any port on any switch can provide full speed access to any other port on the network, subject to bandwidth availability.

FC
See Fibre Channel.

FCP
See Fibre Channel Protocol.

FCP mirror
A Fibre Channel Protocol (FCP) Mirror keeps a consistent image of data on each of two disks. Write commands are duplicated and sent down two paths simultaneously.

fiber
Short for fiber optic cable.

fiber optic cable
A technology that uses glass (or plastic) threads (fibers) to transmit data. A fiber optic cable consists of a bundle of threads, each of which is capable of transmitting messages modulated onto light waves.

fibre
Short for Fibre Channel.

Fibre Channel
A high speed network transmission technology that is capable of transferring data between two ports at up to 100 megabytes/second (higher speeds in the future) over a distance of up to 20 Km. Fibre Channel supports point-to-point, arbitrated loop, and switched topologies. The American National Standards Institute (ANSI), a private, non-profit organization, administers and coordinates the U.S. voluntary standardization and conformity assessment system standards for Fibre Channel.

Fibre Channel Protocol
The mapping of SCSI-3 operations to Fibre Channel.

firmware
The software embedded in a piece of hardware to control that hardware. Generally, firmware can be upgraded and is placed on an electrically erasable programmable read-only memory (EEPROM).

front end
The segment of the hardware configuration that provides input to the system. From the appliance perspective, the front end is from the target HBA(s) up to and including the host(s).

Gbps
Gigabits per second.

Gigabit
For data storage, a Gigabit equals 1024 bits. When referring to data transfer rates, a Gigabit equals 1,000,000,000 bits.

Gigabit Ethernet
An Ethernet protocol running at 1 or 2 Gbps over a fiber medium.

HBA
See Host Bus Adapter.

heartbeat
The network connection that the appliance nodes use exclusively to communicate with each other.

heterogenity
In networking, to establish connectivity between hosts on networks using differing protocols, operating systems, or hardware. To do so may require traversing several networks in between, each of which may be of yet another type. These different networks may be Ethernets, token rings, point-to-point links, or switched networks of various flavors, and each is likely to have its own addressing scheme, media access protocols, service model, and so on. The challenge of heterogeneity is to provide a useful and fairly predictable host-to-host service over this mix of different components.

host
Computing platform connected into the front-end of the SANLink appliance and accessing data from LUNs assigned to it through the appliance. The initiator in a SCSI transaction.

host bus adapter (HBA)
An adapter that provides a Fibre Channel or SCSI interface for a host's data bus.

I/O
See Input/Output.

interoperability
The ability for software and hardware from multiple vendors to communicate.

initiator
A SCSI device or Fibre Channel HBA configured to issue SCSI commands to send to a Target device.

input/output (I/O)
Data traffic

Internet Protocol (IP)
A protocol that specifies the format of packets, also called datagrams, and the addressing scheme. Most networks combine IP with a higher-level protocol called Transport Control Protocol (TCP), which establishes a virtual connection between a destination and a source.

IP
See Internet Protocol.

IP mirror
An Internet Protocol (IP) Mirror keeps a consistent image of data across two disks. Write commands are duplicated and sent down two paths simultaneously.

JBOD
Abbreviation for just a bunch of disks.

journal file
Te journal file tracks new replication requests on the source LUN when a mirror is either paused or the source LUN and target LUN cannot communicate. The journal file enables the LUNs to be quickly and efficiently resynchronized, ensuring the data is consistent.

LAN
See Local Area Network.

latency
The delay in service response. A measurement of time that corresponds to how long it takes a single bit to propagate from one end of a network to another.

link
A connection between two Fibre Channel protocols consisting of a transmit fiber and a receive fiber.

load balancing
Distribution of workload over multiple systems in a cluster.

local
In a network, a topology in which all connections are running on the same subnet of either Fibre Channel Protocol (FCP) or Internet Protocol (IP). In general computing, resources residing on the machine that a person is using. Opposite of remote.

Local Area Network (LAN)
A computer network that spans a relatively small area. Most LANs are confined to a single building or group of buildings. Usually a LAN is less than one kilometer in length.

logical disk
A set of contiguously addressed member disk blocks that is part of a single virtual disk-to-member disk mapping.

Logical Unit Number (LUN)
An identifier assigned to a subunit of storage on a target device created by partitioning and pooling of storage drives. Used in the context of devices connected to a SCSI controller. Each device on a SCSI controller has a SCSI ID, but each SCSI ID may have several LUNs that translate to several partitions, which may be scattered across multiple storage devices.

logical volume
A virtual amount of storage made up of logical disks. Also referred to as a virtual disk, volume set, or partition.

LUN
See Logical Unit Number.

mapping
A term used to describe the process of assigning a storage partition to an application server or host.

Mbps
Short for megabytes per second, a measure of data transfer speed. Mass storage devices are generally measured in Mbps.

mirror
A continuously updated copy of one or more LUNs.

mirroring
Duplicating data from one storage device to another for data backup.

network
A simple network can consist of two or more components directly connected by some physical medium, such as a coaxial cable or an optical fiber. This physical medium is called a *link*, and the components it connects to are *nodes*. A node can be a specialized piece of hardware rather than a computer.

network adapter
The network adapter contains a signaling component that actually encodes bits into signals at the sending node, and decodes signals into bits at the receiving node. Signals travel over a link between two signaling components, and bits flow between network adapters.

network interface card (NIC)
An adapter inside the computer that connects to a shared cable to facilitate physical connection between the computer and the network.

NIC
See network interface card.

node
A node is the processing center of the appliance, including the central processing unit, volatile and non-volatile memory, and supporting hardware.

partitioning
The division of storage space into smaller, independently accessible areas. Used to break a single physical device into one or more smaller virtual devices. Can be used to share storage from a single disk across multiple host operating systems.

point-to-point
Fibre Channel point-to-point is one of three topologies for Fibre Channel systems. Point-to-point architectures consist of two nodes that are directly connected. It allows two N_Ports to be connected via a single link. Larger configurations can be created by providing multiple N_Ports on each node. Each point-to-point circuit provides the full bandwidth supported by the N_Ports.

protocols
Rules established for communicating between devices.

redundant
A fault-tolerant system or component that utilizes an identical system or component that is in stand-by mode, waiting to take over the work of the primary unit in case of failure.

remote
In a network, an Internet Protocol (IP) topology that extends beyond the local subnet. In general computing, resources residing on the machine other than that which a person is using. Opposite of local.

replication
In database management, the ability to keep distributed databases synchronized by routinely copying the entire database or subsets of the database to other servers in the network. There are various replication methods. Primary site replication maintains the master copy of the data in one site and sends read-only copies to the other sites. In a workflow environment, the master copy can move from one site to another. This is called "shared replication" or "transferred ownership replication." In symmetric replication, also called "update-anywhere" or "peer-to-peer replication," each site can receive updates, and all other sites are then updated. Failover replication, or hot backup, maintains an up-to-date copy of the data at a different site for backup.

SANOS
The operating system of the appliance.

scalability
The ability of a system to support growth. In a storage context, a scalable system is capable of efficiently handling the growth from a small system to a large system, in terms of capacity, performance, availability, connectivity, and/or manageability.

SCSI
See Small Computer System Interface.

server
A computer that manages network resources.

Simple Network Management Protocol (SNMP)
A protocol used to poll devices for status information. The management platform is the SNMP manager; the managed device contains an SNMP agent. SNMP works by sending messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

Small Computer System Interface (SCSI)
A parallel interface standard used for peripheral devices, most commonly printers and disk drives. The following varieties of SCSI are currently implemented:

- SCSI-1—Uses an 8-bit bus, and supports data rates of 4 MBps.
- SCSI-2—Same as SCSI-1, but uses a 50-pin connector instead of a 25-pin connector, and supports multiple devices. This is what most people mean when they refer to plain SCSI.
- Wide SCSI—Uses a wider cable (168 cable lines to 68 pins) to support 16-bit transfers.
- Fast SCSI—Uses an 8-bit bus, but doubles the clock rate to support data rates of 10 MBps.
- Fast Wide SCSI—Uses a 16-bit bus and supports data rates of 20 MBps.
- Ultra SCSI—Uses an 8-bit bus, and supports data rates of 20 MBps.
- SCSI-3—Uses a 16-bit bus and supports data rates of 40 MBps. Also called Ultra Wide SCSI.
- Ultra2 SCSI—Uses an 8-bit bus and supports data rates of 40 MBps.
- Wide Ultra2 SCSI—Uses a 16-bit bus and supports data rates of 80 MBps.

SNMP
See Simple Network Management Protocol.

storage area network (SAN)
The implementation of connectivity provided by Fibre Channel between host and storage. This connectivity provides a true network of storage devices and the hosts that access the storage, creating a centralized resource of disk drives and storage subsystems.

storage device
A physical device used for storing data, available in a variety of media (e.g., disks, tapes, optical).

storage pooling
A logical view of multiple storage devices presented as one virtual storage pool to the end-user. Allows for the dynamic management of SCSI and FC back-end devices and access control and allocation mechanisms for storage units.

switch
A network infrastructure component to which multiple nodes attach to exchange data. Unlike hubs, switches maintain full bandwidth capacity to all attached devices. They also are capable of switching node connections from one to another.

synchronous
Synchronous transfer is a mirroring technique in which write commands are duplicated, sent down two paths simultaneously, and confirmation for both writes are received before the acknowledgement is generated.

target
A SCSI device, Fibre Channel HBA, or RAID controller configured to receive SCSI commands.

TCP/IP
See Transmission Control Protocol/Internet Protocol.

throughput
The number of I/O requests satisfied per unit of time, expressed in requests per second.

timeout
Occurs when a function is aborted after a reasonable amount of time for completion has been exceeded.

Transmission Control Protocol (TCP)
Communication protocol used to enable two hosts to establish a connection and exchange data. TCP delivers packets of data in the same order in which they were sent.

Transmission Control Protocol/Internet Protocol (TCP/IP)

A suite of communication protocols, in which layers three and four are combined, used to connect hosts on the Internet.

virtualization

Method by which a storage resource may be used as an asset within a single storage pool, and may be allocated, divided, or represented in a manner different from its physical connection and/or capacity. Virtualization allows for the allocation of storage from the pool to individual hosts. Storage is then presented to attached hosts completely independently of the actual physical configuration of the disks. Virtualization allows storage configuration, data migration, LUN partitioning, and LUN expansion to be possible.

wide area network (WAN)

Larger networks that interconnect geographically distant computers through a public network (such as a telephone system).

zoning

A term used by some switch companies to denote the division of a SAN into subnets that provide different levels of connectivity or addressability between specific hosts and devices on the network. Routing tables are used to control access of hosts to devices. This zoning can be performed by cooperative consent of the hosts or can be enforced at the switch level. In the former case, hosts are responsible for communicating with the switch to determine if they have the right to access a device.

index

A

adding
  display nodes 42
  mirror nodes 156
  peer nodes 46
appliance
  benefits of 16
  installation options 16
  standard features 17
  upgrade options 21
application logs, viewing 32
asynchronous IP mirrors
  defining 160
  mirror states 169
audience 10
authorized reseller, HP 14
Auto Path 22

B

backing up files
  procedure 204

C

cascaded appliances
  defining 19, 123
  relation to cross-appliance mirrors 122
central WAN configuration 153
changing
  user passwords 106
  user roles 101
command line interface 29
command types for managing switches 200
configuring
  IP mirror environment 150, 153
  nodes 45
console
  about 30
  application logs 32
  event details 36
  log filters 39
  log viewer 31
  security logs 33
conventions
  document 11
  equipment symbols 12
  text symbols 11
creating
  expansions 73
  IP mirrors 160
  LUN maps 64
  new roles 99
  new users 104
  partitions 69
  point-in-time images 192
  scripts 209
cross-appliance mirrors
  cascaded appliance configuration 123
  defining 122
  deleting 145
  exporting LUNs 126
  pausing 141
  resynchronizing 143

D

defining
  cascaded appliances 123
  cross-appliance mirrors 122
  direct path for IP mirror 156
  FCP mirrors 110
  IP mirrors 148
  LUNs 62
  physical LUNs 62
  point-in-time images 186
  proxy path for IP mirrors 156
  virtual LUNs 62
deleting
  cross-appliance mirrors 145
  display nodes 44
  expansions 77

FCP mirrors 120
host 60
IP mirrors 182
LUN maps 67
LUNs 94
partitions 72
point-in-time images 195
user roles 103
users 107
direct path for IP mirror 156
display nodes
adding 42
deleting 44
displaying LUN information 78
document
conventions 11
related documentation 10
dual WAN configuration 150

E

editing host information
using sanreg utility 55
using UI or CLI 58
editing role privileges 101
equipment symbols 12
event details, viewing 36
executing a script 211
expansions
creating 73
deleting 77
expanding and merging LUNs 73
exporting LUNs 126

F

FCP mirrors
defining 110
deleting 120
mirror states 114
pausing 116
Fibre Channel mode for appliance target ports 50
files
backing up 204
restoring 205
filters
for logs 39

G

getting help 14

H

help, obtaining 14
hosts
about host connections 54
deleting 60
editing hosts
using sanreg 55
using UI or CLI 58
HP
authorized reseller 14
storage website 14
technical support 14

I

IP mirrors
adding a mirror node 156
asynchronous mirrors 160
creating 160
defining 148
deleting 182
mirror relationships 149
central WAN configuration 153
direct path 156
dual WAN configuration 150
proxy path 156
mirror states
asynchronous 169
synchronous 166
pausing 174
removing a mirror node 183
synchronous mirrors 160

L

logs
application logs 32
log filters 39
log viewer 31
security logs 33
LUN maps
creating 64
defining 63
deleting 67
LUN resizing 68
expansions
creating 73
deleting 77
partitions
creating 69
deleting 72
LUN tree
viewing physical LUNs 82
viewing virtual LUNs 85
LUNs
defining 62
deleting 94
displaying information 78
exporting for cross-appliance mirrors 126

new and changed LUNs 95
performance 90
physical 62
setting location 87
unmapping 67
virtual 62

M

managing switches
command types 200
registering IP addresses 198
mapping LUNs 63
mirror states
asynchronous IP mirrors 169
FCP mirrors 114
synchronous IP mirrors 166
mirrors
cross-appliance
defining 122
deleting 145
exporting LUNs 126
pausing 141
resynchronizing 143
FCP
defining 110
deleting 120
pausing 116
IP
asynchronous mirroring 160
creating 160
defining 148
deleting 182
mirror nodes 156, 183
pausing 174
synchronous mirroring 160
modifying user information
passwords 106
privileges 101
roles 101

N

nodes
display nodes
adding 42
deleting 44
mirror nodes
adding 156
removing 183
node configuration 45
peer nodes 17
adding 46
removing 47

P

partitions
creating 69
deleting 72
rules for 69, 74
passwords, changing 106
pausing
cross-appliance mirrors 141
FCP mirrors 116
IP mirrors 174
peer nodes
adding 46
removing 47
performance of LUNs 90
physical LUNs 62, 82
point-in-time images
common uses 22
creating 192
defining 22, 186
deleting 195
viewing available targets 194
ports, target ports on the appliance 50
privileges
defining 98
editing 101
proxy path for IP mirrors 156

Q

queue drainer
starting the queue 161
stopping the queue 161

R

rack stability, warning 13
related documentation 10
removing
mirror nodes 183
peer nodes 47
rescanning for new/changed LUNs 95
resizing LUNs 68
restoring files
procedure 205
resynchronizing
cross-appliance mirrors 143
roles
adding privileges 99
creating 99
defining 98
deleting 103
modifying 101

S

sanreg utility, for editing host information 55

scripts
creating 209
executing 211
security logs, viewing 33
setting LUN location 87
statistics, LUNs 90
symbols in text 11
symbols on equipment 12
synchronous
FCP mirror states 114
IP mirror states 166

T

target ports
defining for the appliance 50
Fibre Channel mode 50
technical support, HP 14
text symbols 11

U

unmapping LUNs 67
user interface
login 24
navigation 26
users
changing passwords 106
creating new roles 99
creating new users 104
defining roles and privileges 98
deleting roles 103
deleting users 107
editing privileges 101
modifying roles 101

V

viewing
available targets for point-in-time images 194
node configuration 45
performance of LUNs 90
virtual LUNs 62, 85
virtualization
benefits of 18
defining 18

W

warning
rack stability 13
symbols on equipment 12
websites
HP storage 14